



HAL
open science

Sampling, interpolation and detection. Applications in satellite imaging.

Andrés Almansa

► **To cite this version:**

Andrés Almansa. Sampling, interpolation and detection. Applications in satellite imaging.. Signal and Image processing. École normale supérieure de Cachan - ENS Cachan, 2002. English. <NNT : >. <tel-00665725>

HAL Id: tel-00665725

<https://theses.hal.science/tel-00665725v1>

Submitted on 2 Feb 2012

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

École Normale Supérieure de Cachan

THÈSE

présentée par

Andrés ALMANSA

pour obtenir le grade de

DOCTEUR DE L'ÉCOLE NORMALE SUPÉRIEURE DE CACHAN

Spécialité: **Mathématiques Appliquées**

Échantillonnage, interpolation et détection.

Applications en imagerie satellitaire.

Soutenue le 9 décembre 2002 devant le jury composé de :

M.	Luis ALVAREZ	
M.	Vicent CASELLES	<i>Rapporteur</i>
Mme.	Françoise DIBOS	
M.	Jean-Michel MOREL	<i>Directeur</i>
M.	Yves MEYER	
M.	Gregory RANDALL	<i>Rapporteur</i>
M.	Bernard ROUGÉ	<i>Directeur</i>

À Meche

À mes parents et ma soeur

Remerciements

Yves Meyer, Luis Alvarez et Françoise Dibos m'ont fait l'honneur d'accepter d'être membres du jury, et Vicent Caselles, Gregory Randall et Lenny Rudin d'être rapporteurs de cette thèse. Je leur suis très sincèrement reconnaissant pour leur disponibilité ainsi que pour leurs commentaires et questions, grâce auxquelles la version finale de cette thèse a été beaucoup améliorée. Un jury et des rapporteurs si nombreux et si prestigieux a représenté à la fois un grand honneur et un grand défi.

Je tiens tout d'abord à exprimer mes plus vifs remerciements à Jean Michel Morel et Bernard Rougé, pour m'avoir fait si grande confiance en m'offrant la possibilité de travailler sur des sujets de recherche si riches et inspirateurs, et dans un cadre humain et professionnel si exceptionnel. Bernard Rougé m'a fait goûter la beauté de la théorie de Shannon, et il a été une source d'inspiration constante. Je ne sais pas comment lui remercier pour les fréquentes et longues heures de discussion et travail commun qu'il m'a accordé à Cachan et à Toulouse, malgré notre distance géographique. J'ai toujours admiré le grand enthousiasme et l'esprit d'équipe que Jean-Michel Morel a su transmettre à ses élèves et collaborateurs, et non moins ai-je apprécié sa grande disponibilité, et les innombrables et brillantes idées qu'il m'a transmises, avec la plus grande humilité, sans même vouloir en rester le coauteur. Son souci pour mon avenir après la thèse, ainsi que les multiples relations nationales et internationales desquelles il m'a donné l'opportunité de faire partie, sont aussi un signe de sa grande générosité.

Dans ce cadre exceptionnel de travail en équipe, j'ai naturellement eu l'occasion d'avoir des discussions très éclairantes avec virtuellement tous les membres du CMLA et ses visiteurs, et ils méritent tous mon plus sincère remerciement. Je voudrais simplement souligner dans les lignes qui suivent quelques contributions particulièrement importantes que j'ai eu la grâce de recevoir.

Tout mon travail sur l'échantillonnage régulier et la restauration, contenu dans les premiers deux chapitres a été le fruit d'une réflexion très intense, dès le début de ma thèse avec Sylvain Durand et Bernard Rougé. Mes travaux sur l'interpolation de modèles de terrain ont bénéficié de la remarquable collaboration de Frédéric Cao dans

les aspects EDP, et de la maîtrise de Yann Gousseau des aspects statistiques, ainsi que d'un souci commun de trouver la meilleure solution au problème. Mon travail sur l'échantillonnage irrégulier n'a aboutit que grâce au soutien initial de Stéphane Jaffard, et le logiciel et commentaires généreusement transmis par Thomas Strohmer, mais surtout grâce à l'insistance et persévérance de Bernard Rougé dans la recherche d'une réponse à ses conjectures. Finalement, Jean-Michel Morel, Agnès Desolneux et Lionel Moisan m'ont introduit dans ce monde fascinant de la théorie de la Gestalt qui a conduit à mon travail sur la détection de points de fuite, travail dont une grande partie a été inspirée par l'excellent travail de Sébastien Vamech, et les discussions qui ont suivi sa participation au groupe de travail sur la géométrie intégrale.

Je souhaite également remercier très chaleureusement Vicent Caselles et Andreu Solé, pour leur critique constructive à mon travail sur l'interpolation par EDPs, ainsi que leur aide éclairante avec François Malgouyres et Lionel Moisan sur la restauration par variation totale et leur algorithmes de minimisation. Mon travail sur l'interpolation a été aussi énormément facilité par les données de Carlos López, et des discussions avec Pierre Soille et Corinne Vachier sur les distances géodésiques. Dans l'échantillonnage irrégulier j'ai eu l'honneur de recevoir des très précieuses suggestions de la part de Yves Meyer et de Patrick-Luis Combettes. Finalement les commentaires et suggestions de Pascal Monasse, ainsi que la soigneuse lecture des rapporteurs anonymes de IEEE PAMI, ont contribué à améliorer nettement le dernier chapitre sur la détection de points de fuite.

Lenny Rudin deserves a very special acknowledgement. I very much enjoyed my visit to his company, and the opportunity to contribute to his excellent video-processing software. My work on vanishing point detection was mainly motivated by the photogrammetric applications at Cognitech and is due to a large extent to Lenny Rudin's ideas and encouragement. The many discussions he granted me on sampling, restoration, and interpolation almost certainly left their traces in this thesis.

Les conditions de travail au CMLA ont été excellentes. J'en remercie tous les membres du laboratoire, et leurs équipes d'administration du système, secrétariat et bibliothèque, ainsi qu'aux doctorants de la salle informatique au rez-de-chaussée, et tous ceux qui ont contribué à la création d'une atmosphère de travail très conviviale.

Plusieurs personnes ont joué un rôle très important dans mon parcours scientifique avant d'arriver au CMLA. I would like to specially thank Tony Lindeberg for he introduced me in the field of image processing, for the many things he taught me during my master's thesis, and for his generous help during my transition from MSc to PhD stu-

dies. Roberto Oliveira-Mattos, Rosario Curbelo y Gustavo Drets merecen mi más sincero agradecimiento por el excelente grupo de trabajo que crearon en Uruguay, y por los vínculos y medios materiales que me permitieron continuar mis estudios. Cuando debí dejar dicho grupo para comenzar mi tesis, Gregory Randall, Alicia Fernández, Alvaro Gomez, Ana Spangenberg, Sebastian Sasías, Sergio Nemachnov, y muchos otros colegas del CeCal y del IIE asumieron generosa y rápidamente muchas de las responsabilidades que poco a poco deje atrás, facilitando así el comienzo de mi tesis. Gregory Randall me ha considerado como un miembro más de su equipo, a la hora de buscar el mejor lugar para mis estudios de doctorado, así como a la hora de hacerme un lugar al regreso. Por su amistad, y por el entusiasmo y profesionalismo que él ha sabido contagiar a sus colaboradores, quiero expresarle mi más sincero agradecimiento.

Dès notre arrivée en France, nous avons été chaleureusement accueillis par nos amis Nouredine, Raphaël, Aurélie, Philippe et Kyung Lan. No menor agradecimiento se merecen mis compatriotas Richard y Claudia, Gianfranco, Ana, Pablo, Nico, Andrea, Denis y Luciana, Laura y Laurent, Gonzalo, Miguel y Rosario, Gustavo, y los casi compatriotas Hilda, Claudia y Guillaume, Vincent y Romina, Pierre, Alejandro por su apoyo constante y por hacernos sentir casi como en casa, durante estos años de tesis. Wie zu Hause haben wir uns auch gefühlt während unser Besuch der Familien Köhncke und Nölting in Mainz, Dreieich, Hamburg, und Cabrils. Sie sind unsere zweite Familie geworden, während unser Aufenthalt in Europa.

Il me tient particulièrement à coeur remercier à chacun des membres de ma communauté CVX à Paris, ainsi qu'à Jean-Marc Fournon, Marika Thomas, Alain Thomasset, Christian Mellon, Philippe Deterre, le groupe Transgenèse, et la communauté de l'Arche à Paris. Non seulement leur écoute et présence tout au long de ces années en France, m'ont permis de croître personnellement. Ils ont aussi su me donner le soutien nécessaire aux moments les plus difficiles de cette thèse. Quisiera hacer extensivo este agradecimiento a mi hermana Mónica, a Alejandro, Marcelo, Rafa y Gabriel, quienes de cerca o a la distancia me han sabido apoyar cuando mas los necesitaba.

Por último, quiero expresar mi más profundo agradecimiento a mi familia y mi esposa. A mis padres, mi hermana Mónica, mis suegros Pilar y Carlos y mi cuñada Magda, por la dignidad y generosidad con que aceptaron mi ausencia y la de mi esposa, y la pesada carga que les dejamos durante estos años en Francia. Igual agradecimiento quiero expresar a mis padres, Susanne y José Antonio, cuyo esfuerzo y determinación de toda una vida me permitieron obtener la educación y los medios que condujeron a esta tesis. Y a mi querida Meche quiero dedicarle mi más sincero agrade-

cimiento, no solo por su paciencia y comprensión, sino por su apertura a esta aventura que nos abrió a ambos nuevos horizontes, y por tantas otras cosas que no cabrían en estas páginas.

Résumé


Cette thèse aborde quelques-uns des problèmes qui surviennent dans la conception d'un système complet de vision par ordinateur : de l'échantillonnage à la détection de structures et l'interprétation. La motivation principale pour traiter ces problèmes a été fournie par le CNES et la conception du satellite SPOT5 d'observation terrestre (qui a été lancé en avril cette année), ainsi que par les applications de photogrammétrie et vidéo-surveillance chez Cognitech, Inc. pendant les étapes finales de ce travail. Au-delà de ces motivations initiales, la plupart des sujets sont traités avec une généralité suffisante pour être d'intérêt pour d'autres systèmes de vision par ordinateur, que ce soient des appareils photo numériques, ou des systèmes de surveillance semi-automatiques.

Si l'on fait une abstraction très grossière, la chaîne de traitement d'image d'un satellite d'observation terrestre est composée des parties suivantes :

À bord : Un système optique reçoit la lumière du paysage observé, et le focalise sur une matrice de capteurs, qui mesure la luminance moyenne reçue par chaque capteur. Ce signal analogique est envoyé à un convertisseur analogique/numérique, et le signal numérique résultant est débruité et compressé pour sa transmission à la base terrestre.

Segment sol : Le signal discret reçu est décompressé et restauré pour éliminer le flou et les artefacts de compression. Certaines corrections géométriques sont parfois réalisées afin de compenser les déformations dues à une prise de vue pas parfaitement verticale, ou bien pour normaliser la géométrie épipolaire d'une paire stéréo. Des corrections plus fines encore sont nécessaires afin de compenser les micro-vibrations subies par le satellite pendant l'acquisition de l'image.

Applications : Les images stockées peuvent être ultérieurement utilisées pour une myriade d'applications qui essayeront de prendre des mesures sur l'image. Ces mesures devraient être les plus précises possibles, ou bien les erreurs devraient être pour le moins prédictibles ou avoir des statistiques connues. Ceci impose des contraintes



spécifiques, bien différentes de celles imposées par des systèmes visant un consommateur final, auquel cas des critères d'apparence visuelle et d'esthétique sont plus critiques. Une des applications principales de l'imagerie satellitaire (et en particulier de SPOT5 qui contient un instrument spécialement conçu à cette fin) est la construction des modèles numériques d'élévation et de terrain à partir de paires stéréo. Récemment, l'utilisation de paires stéréo quasi-simultanées avec une petite différence d'angle de prise de vue est en train d'être sérieusement considérée. L'avantage d'une telle disposition, est qu'en évitant au maximum les occlusions et les changements d'illumination d'une prise de vue à l'autre, elle permet d'obtenir une carte d'élévation assez complète par corrélation de la paire d'images. En revanche, la magnitude des vecteurs dans la carte de disparité est beaucoup plus petite qu'en photogrammétrie traditionnelle, et en conséquence beaucoup plus sensible aux erreurs dans les mesures prises sur l'image et son interpolation. Voici l'une des raisons pour lesquelles ce projet a mis l'accent sur la conception du système d'acquisition d'image, et sur les procédures à utiliser pour échantillonner et interpoler correctement les images produites par ledit système. Cet effort de précision s'est montré très rentable quand nous nous sommes attaqués à d'autres problèmes de détection de structures.

Dans ce cadre, nous avons concentré notre travail sur cinq questions.

1. Comment choisir le "meilleur" système d'acquisition parmi une série de possibles ?
2. Comment quantifier et éviter l'effet de repliement spectral (aliasage) dû à l'échan-

tillonnage ?

3. Comment corriger l'échantillonnage irrégulier dû aux micro-vibrations du satellite ?
4. Comment remplir les régions dans lesquelles une paire stéréo n'a pas réussi à fournir une estimation fiable de l'élévation du terrain ?
5. Comment étendre les méthodes de corrélation de paires stéréo à des zones urbaines, dans lesquelles la carte d'élévation est bien moins régulière que dans le cas de modèles de terrain ?

Cette thèse n'a pas l'ambition de fournir des réponses complètes à toutes ces questions ; nous avons tout simplement essayé de contribuer à la compréhension de chaque problème. Notons, que la première et la dernière question sont spécialement générales. Notre réponse sera donc dans ces deux cas, particulièrement partielle.

La première question est assez liée à la seconde et à la troisième. On peut, bien sûr, étudier les propriétés et les possibilités offertes par chacun des réseaux d'échantillonnage, comme nous le faisons dans l'appendice A, mais pour obtenir une réponse plus complète à cette question nous devons estimer dans quelle mesure chacun des systèmes permet d'éviter les artefacts les plus communs pendant l'acquisition de l'image, tels que le flou, le bruit, le repliement spectral, et les perturbations du réseau d'échantillonnage (ou "jitter", comme ce phénomène est nommé dans une partie de la littérature anglophone sur ce sujet). Ces questions sont traitées dans le chapitre 1, en ce qui concerne le repliement spectral principalement, mais aussi le bruit et le flou, dans une certaine mesure, et ensuite dans le chapitre 2, en ce qui concerne l'échantillonnage perturbé. Notre *Leitmotiv* dans ces deux chapitres (ainsi que l'appendice A et l'annexe B) est le souci de répondre à la première question de différents points de vue. Ces travaux mettent au centre de nos préoccupations la théorie de Shannon de l'échantillonnage et la théorie de l'information, et ils utilisent les outils de l'analyse de Fourier harmonique et non-harmonique. Ils ont mérité donc d'être regroupés dans la partie I de cette thèse, sur l'échantillonnage d'images et l'interpolation de Shannon.

Quand il s'agit de combler des trous dans un modèle d'élévation ou de terrain (question 4 et chapitre 3), nous devons adopter un modèle d'interpolation différent. Le modèle de Shannon, basé sur la supposition que l'image est essentiellement à bande limitée, n'est plus valable car : (i) la surface du terrain n'a aucune raison d'être à bande limitée ; (ii) les mesures d'altitude que nous faisons ne sont pas pré-filtrées non plus par une " fonction transfert " à bande limitée ; (iii) même si le modèle de Shannon était valable pour représenter la surface du terrain, il y a un problème d'échelle, car la taille

des trous à remplir est souvent beaucoup plus grande que le taux d'échantillonnage associé à toute limite raisonnable pour la bande spectrale. Pour cette raison, le modèle de Shannon produit soit des interpolateurs beaucoup trop réguliers qui ne peuvent pas satisfaire les conditions aux bords, soit des interpolateurs beaucoup trop oscillants. Aussi, dans le chapitre 3, nous nous intéressons à des modèles plus géométriques, basés sur des équations aux dérivées partielles, qui permettent de mieux exprimer la structure géométrique particulière des surfaces de terrain.

Enfin, nous arrivons à la dernière question, sur l'extension des méthodes de corrélation pour les paires stéréo, au cas des paysages urbains. Paradoxalement, bien que les bords contrastés soient essentiels pour permettre aux méthodes de corrélation de fournir des cartes d'élévation précises, la corrélation échoue systématiquement près de ces bords (les arêtes d'un bâtiment par exemple) à cause du phénomène d'adhérence décrit par Delon et Rougé [57]. En conséquence, les méthodes de corrélation ne fournissent pas des estimations fiables de la hauteur dans les régions proches d'un bord très contrasté, et si l'on utilise des méthodes d'interpolation comme ceux du chapitre 3, les murs et les coins d'un immeuble seraient lissés, arrondis, et d'autres artefacts apparaîtraient. Il semble assez évident que pour compléter ces limitations des méthodes de corrélation une meilleure approche consisterait à détecter et localiser précisément les bords et d'autres caractéristiques, et à faire ensuite une mise en correspondance entre les deux images. Plus ces caractéristiques seront groupées pour former des structures géométriques complexes, moins il sera probable de les trouver par hasard, et la mise en correspondance sera plus significative. Ceci nous amène au programme de recherches de la théorie de la Gestalt de Wertheimer [173], Metzger et Kanizsa [108], et leur implémentation computationnelle introduite par Desolneux, Moisan et Morel [62, 63] basée sur le principe de Helmholtz et les modèles *à contrario*.

Dans le chapitre 4, nous suivons ce programme, et nous nous concentrons sur deux Gestalts omniprésentes dans les scènes urbaines, qui sont particulièrement utiles pour la création de modèles d'élévation à partir de paires stéréo, ou même des modèles architecturaux à partir d'une seule prise de vue : les segments de ligne droite et leurs points de fuite. L'utilisation des alignements et points de fuite dépasse largement le cadre de la conception de systèmes d'imagerie satellitaire. Ils peuvent être utilisés pour calculer la géométrie épipolaire, calibrer une partie des paramètres internes ou externes de la caméra ou même prendre des mesures très précises des dimensions tridimensionnelles d'un objet, à partir d'une seule ou plusieurs prises de vue de l'objet, et ceci en utilisant des connaissances minimales sur la scène, telles que l'orthogonalité

ou l'appartenance à un même plan. De fait, cette dernière application a été proposée par Cognitech, Inc., dans le but d'aider à la résolution de certains problèmes de photogrammétrie, dans lesquels l'entreprise est spécialisée, et celle-ci a été la motivation principale pour nous intéresser à ce problème.

Guide de lecture Chacun des chapitres 1 à 4 ainsi que l'appendice A constitue en lui-même une unité autonome, et peut être lu indépendamment des autres. La seule exception est l'annexe B, qui contient à la fois des matériaux permettant de mieux comprendre les contenus et certains des choix de la partie I, ainsi que des conclusions d'ensemble de la partie I. Une progression assez logique de lecture pourrait être la suivante : appendice A, partie I, annexe B, chapitre 3, chapitre 4, sachant qu'il serait préférable de lire certaines parties de l'annexe B là où elles sont citées dans la partie I. Cette progression n'a pas été respectée pour des raisons de confidentialité dans le cas de l'annexe B, et pour éviter de donner trop d'importance à l'appendice A, qui contient, contrairement aux chapitres 1 à 4, très peu de résultats originaux. Il s'agit simplement d'une description de l'état de l'art, dont une partie a motivé le reste du travail. De plus, certains aspects de cette partie sont repris plus rigoureusement, et avec plus de détails dans le chapitre 1.

Table des matières

Remerciements	v
Résumé	ix
Table des matières	xv
Overview	xix
Notations	xxix
I Image Sampling and Shannon Interpolation	1
1 Regular Sampling, Resolution and Restoration	3
1 Introduction	4
2 Fourier Analysis in non-square grids	9
2.1 Sampling grids and tilings of the plane	9
2.2 Shannon's sampling theorem.	11
2.3 Choice of the reciprocal cell.	13
2.4 Critical Sampling.	15
3 Modeling of image acquisition systems	17
4 Resolution and information measures	25
4.1 Nominal Resolution	25
4.2 Effective resolution	27
4.3 Information density.	30
5 Optimal Reciprocal Cell	33
6 Examples and Applications	38
6.1 Optimal Reciprocal Cells	38
6.2 Resolution and Information measures	38

6.3	Application to Restoration	44
7	Discussion and future work	53
A	Proof of Shannon's sampling theorem	55
A.1	Prerequisites	56
A.2	Shannon's sampling theorem	59
A.3	Suitability of the hypotheses in real systems	63
B	Proof of information bound	64
C	Sufficient conditions for admissibility	67
2	Irregular Sampling	69
1	Introduction	69
2	Existence theorems	71
2.1	Problem statement.	71
2.2	Perturbed critical sampling	72
2.3	Irregular over-critical sampling	74
3	Available Reconstruction Algorithms	77
3.1	Based on Kadec's condition (perturbed sampling)	77
3.2	Based on Beurling-Landau type conditions (irregular sampling)	77
4	A new reconstruction algorithm	84
4.1	Convergence analysis	85
4.2	Numerical approximation	89
5	Experiments	89
6	Conclusions and perspectives	96
A	Proofs of convergence analysis.	97
II	Non-Shannon Interpolation	105
3	Interpolation of Digital Terrain Models	105
1	Introduction	106
2	The AMLE model	107
3	Relations to Previous Work on DEM interpolation	111
3.1	Geodesic distance transformations.	111
3.2	The thin plate model	112
3.3	Kriging	115
4	Experiments	117
4.1	A simple combination of the AMLE and thin-plate models	118

5	Discussion and Further Work	119
A	AMLE and weighted geodesic distance interpolations	133
	A.1 Review of weighted geodesic distance interpolation	133
	A.2 Equivalence with AMLE in the twice differentiable case	134
	A.3 Partial results in the weakly differentiable case	138
	A.4 Experiments	146
B	Numerical Analysis	147
	B.1 AMLE model using finite differences	148
	B.2 Thin-plate using finite differences	149
C	Algorithmic complexity	152
	C.1 AMLE using finite differences	153
	C.2 AMLE through iterated geodesic distance transforms	153
	C.3 Thin-plate and Kriging through radial basis functions	153
	C.4 Thin-plate through finite differences	154
III Feature Detection		155
4	Vanishing point detection	157
1	Introduction	158
2	Detection of Line Segments	160
	2.1 Helmholtz principle and alignment detection	160
	2.2 Proposed refinements	162
3	Detection of Vanishing Points	164
	3.1 Meaningful vanishing regions	166
	3.2 Probability of a line meeting a vanishing region.	167
	3.3 Partition of the image plane into vanishing regions	171
	3.4 Final remarks	176
	3.5 Algorithm	179
4	Experiments and Discussion	184
5	Conclusions and Perspectives	185
A	Proof of the main results on integral geometry.	197
Appendix		201
A	Review of Hexagonal Sampling ¹	203

1	Introduction	203
2	Digitalization	204
2.1	Tilings of the plane	205
2.2	Regular tilings and tilings by regular polygons	208
2.3	Tiling-grid duality	209
2.4	Conclusion	210
3	Properties of Square and Hexagonal Grids	210
3.1	Discrete Connectivity	211
3.2	Optimal Sampling	214
3.3	Recursivity	221
3.4	Miscellaneous properties	222
4	Morphological and level-set operations	223
4.1	Thinning	224
4.2	Distance Transforms	224
5	Linear Filtering	225
5.1	Review of Fourier Analysis	225
5.2	Discrete Fourier transform	227
5.3	Separable Filtering	230
5.4	Conversion between Square and Hexagonal	231
6	Wavelets	234
6.1	Classical construction	234
6.2	Orthogonal wavelet bases with hexagonal symmetry	234
6.3	Hexagonal subband coding	235
6.4	Biorthogonal hexagonal wavelets	237
6.5	Wavelets on irregular grids	237
7	Miscellaneous Operations and Applications	239
7.1	Computer Graphics primitives	239
7.2	Pattern recognition	239
7.3	Operating on Hexagonal Grids	239
7.4	Hardware support for parallel hexagonal image processing	241
8	Summary and Discussion	241
A	Grid conversion algorithm	242

Bibliography **245**

Annexe B. Systèmes d'échantillonnage du CNES **261**

Overview

This thesis deals with a few of the many issues that arise in the design of a complete computer vision system, from sampling and interpolation, to feature detection and interpretation. The main motivation for addressing these topics was provided by the design of the SPOT 5 earth observation satellite launched in April this year, the planning of its successor satellites, as well as a photogrammetry and video-surveillance application at Cognitech, Inc. during the final stages of this work. Beyond these initial motivations, most subjects are treated with full generality and can be of interest for other computer vision systems, let it be *e.g.* digital cameras, or semi-automated surveillance systems.

If we make a very coarse abstraction, the image processing chain of an earth observation satellite consists of the following parts:

On board segment. An optical system gathers the light from the observed landscape, and focuses it at a sensor array, which measures the average luminance received by each sensor. This analog signal is sent to a D/A converter, and the corresponding digital signal is de-noised and compressed for transmission to the earth base segment.

Earth base segment. The received discrete image is decompressed, and restored to eliminate blur and compression artifacts. Geometric corrections are usually performed, in order to compensate for the viewing angle, or to standardize the epipolar geometry of stereo pairs. Even finer corrections are needed to compensate for micro-vibrations of the satellite during acquisition.

Applications. The stored images can later be used for a myriad of applications, which will try to perform measurements on the image which should be as precise as possible or, at least, these errors should be predictable or have known statistics. This imposes specific constraints on the final image, different from consumer applications where the visual appearance and aesthetic criteria are critical features. One of the main appli-

cations of satellite imaging (and in particular of the SPOT5 satellite which contains a special instrument for this purpose) is the construction of digital elevations models from stereo pairs. Recently, almost simultaneous stereo pairs with an unusually small difference in viewing angle have been seriously considered. The advantage of this disposition, is that by avoiding changes and occlusions, it can provide a quite complete height computation by correlation of the image pairs. On the down side, disparity maps are much smaller than in traditional photogrammetry, and therefore more sensitive to errors in image measurements and interpolation. This is one of the reasons why a special emphasis was put throughout this project on the design of the acquisition system, and on how to correctly sample and interpolate images. This effort paid off later, when we addressed apparently unrelated feature detection issues.

Within this framework we concentrated our work on five questions

1. How to choose the “best” acquisition system among a series of possibilities?
2. How to quantify and avoid the aliasing effect due to sampling?
3. How to compensate for the irregular sampling due to micro-vibrations of the satellite?
4. How to fill-in the regions where the correlation of stereo pairs failed to provide a height value?
5. How to extend the stereo-pair correlation methods to urban areas, where the height map is much less regular than in terrain models?

This thesis does not have the ambition to provide complete and final answers to all of these questions ; we simply attempt to provide some more insight on each subject. Note that the first and last questions are especially general and our answer will be particularly partial.

The first question is quite related to the second and the third ones. We can study the properties and possibilities of the different possible sampling grids as we do in appendix A, but a more complete answer to this question involves comparing to what extent each of the systems allows to avoid common artifacts during image acquisition such as blurring, noise, aliasing, and perturbations of the sampling grid (or sampling “jitter” as it is referred to in part of the literature). These questions are addressed in chapter 1, concerning mainly aliasing, and also to some extent noise and blurring, and in chapter 2, concerning jitter. Our Leitmotiv in these two chapters (together with appendix A and annex B) is to try to answer the first question from different points of

view. Thus, they all share a focus on Shannon’s sampling and information theory and they involve harmonic and non-harmonic Fourier analysis. We grouped these three chapters into part I of this thesis, on image sampling and Shannon interpolation.

When it comes to filling in holes in digital elevation models (question 4 and chapter 3), we have to switch to a different interpolation model. The Shannon model, based on the assumption that the image is essentially band-limited, is no longer valid because (i) the terrain surface has no reason to be band-limited; (ii) the height measurements we make are not pre-filtered by a band-limited “transfer function” either; (iii) even if the Shannon model was valid to represent terrain surfaces, there is a scale problem, since the size of the holes that need to be filled is often much larger than the sampling rate corresponding to any reasonable band limit. Therefore the Shannon model leads to either much too smooth interpolators which cannot satisfy the boundary conditions, or to much too oscillating interpolators.

For this reason in chapter 3 we turn to more geometric models based on partial differential equations, that allow to better express the particular geometric structure of terrain surfaces.

Finally, we come to the last question, on how to generalize stereo pair correlation methods to the case of urban landscapes. Paradoxically, even if contrasted edges are essential for correlation methods to provide accurate elevation maps, they systematically fail close to the edges (of buildings for instance) due to the adherence effect described by Delon and Rougé [57]. Therefore correlation methods do not provide height measurements close to the edges, and if we use interpolation methods such as those in chapter 3, walls and corners of buildings become smoothed out, and other artifacts are created. It seems quite obvious that to complement these limitations of correlation methods, a better approach would be to detect and precisely locate edges and other features, and then match them between both images. The more these features are grouped into complex geometric structures the less they will be likely to appear, and the more a match (between two similar such complex structures in both images) will be significant. This leads us to the full Gestalt programme of Wertheimer [173], Metzger, and Kanizsa [108], and their computational implementation based on the Helmholtz principle and a contrario models introduced by Desolneux et al. [62, 63].

In chapter 4 we follow this programme, concentrating on two features that are particularly common in urban scenes, and useful for the creation of elevation models from stereo pairs or even architectural models from single views: line segments and their vanishing points. The use of alignments and vanishing points goes beyond the design

of a satellite imaging system. They can be used to compute the epipolar geometry, calibrate part of the the internal or external parameters of a camera, or even take precise 3D measurements of length ratios from a single or several views of the scene, using a minimal a priori knowledge of the scene like coplanarity or orthogonality relations. Actually, this last application was proposed by Cognitech Inc., in order to help solve some of the photogrammetry problems in which the company is specialized, and this provided the main motivation for working in this area.

Reading guide. Each of the chapters 1 to 4, as well as appendix A makes up a self-contained unit, and can be read in isolation of the other chapters. The only exception is annex B, which contains *(i)* supporting materials which may help better understand the contents and certain choices made in part I, and *(ii)* overall conclusions concerning part I. A logical reading sequence could be: appendix A, part I, annex B, chapter 3, chapter 4, taking into account that it could be preferable to read certain parts of annex B at the point where they are cited in part I. This sequence was not respected in the organization of the thesis for confidentiality reasons in the case of annex B, and to avoid giving too much importance to appendix A, which is mainly an overall review of the state of the art and contains much less original material than the rest of the chapters. The parts of appendix A that are important for the rest of the thesis, are revisited in more detail in chapter 1, so it may well be skipped by the reader who wishes to go directly to the main contributions of this thesis.

In the following sections we summarize the main contributions of each of the four chapters and the appendix.

Review on Hexagonal Sampling

In this appendix we give a brief account of the different possible image sampling grids and their properties. The main reason for conducting this study, is that recent developments at the French Space Agency (CNES) allow for the realization of high resolution satellite acquisition systems that do not need to locate the sensors on a traditional square grid. We start with an analysis of the geometric constraints on symmetry and regularity that leave only two possibilities in the two-dimensional case, namely: square and hexagonal grids. Next, we review the main properties of these two grids, concluding that: *(i)* the hexagonal grid requires a smaller sampling density to represent the same amount of information, under isotropy hypothesis; *(ii)* discrete topology (and

hence morphological operators) are more conveniently expressed in terms of a single connectivity concept, whereas square grids require two complementary concepts; *(iii)* the hexagonal grid allows for the construction of wavelet families with better orientation selectivity, more accurate distance transforms, and isotropic separable filters with a slightly larger computational effort; *(iv)* the discrete Fourier transform can be computed on an hexagonal grid using the same FFT algorithm as in the square grid, if we only accept to periodize the image along its main axes, instead of periodizing it along orthogonal axes as in the square grid. Based on these ideas we propose a new algorithm for converting between square and hexagonal grids using Shannon interpolation, and separable sinc filters.

Optimal reciprocal cell, and applications to restoration and zoom

In view of the many good properties of hexagonal grids, it was decided to study in more detail the resolution of both square and hexagonal sampling systems. An important prerequisite for this study is a thorough development of Fourier analysis and sampling theory in non-orthogonal grids. Even though the main results are available in different places it was difficult to find a single reference containing the development of the results we needed in a unified manner. For this reason we summarize these results in section 5.1 following mainly Morel and Ladjal [135], Bony [36], and Gasquet and Witomski [80].

In this context, following Rougé [147] a sampling system is assumed to be composed of the triplet (Γ, H, N) where Γ represents the sampling grid, H its transfer function, and N the noise statistics, more precisely its standard deviation at each frequency. In addition some knowledge of natural image statistics is assumed, namely the standard deviation F at each frequency. From this information we compute a signal-to-noise ratio

$$\frac{1}{b^2(\xi)} = \frac{|HF|^2(\xi)}{N^2(\xi)}$$

and a signal-to-aliasing ratio

$$\frac{1}{a^2(\xi)} = \frac{|HF|^2(\xi)}{\sum_{\gamma^* \in \Gamma^* \setminus \{0\}} |HF|^2(\xi + \gamma^*)}$$

at each frequency, which allows us to determine the spectral domain (reciprocal cell), which satisfies a minimum specified threshold for both values. The size of this optimal spectral domain, or a weighted average within this domain, provide measures of

effective resolution in terms of equivalent samples per unit area of an ideal system. Alternatively, from these two ratios we can measure the mutual information between a Fourier coefficient of the sampled image and the corresponding Fourier coefficient of the analog image before sampling. When we average this mutual information over the whole frequency domain, we obtain a resolution measure in terms of bits per unit area or bits per sample, where the number of bits measures the portion of the information that is common between the original and sampled image, and not due to noise or aliasing. The maximization of this second resolution measure also leads to an optimal reciprocal cell. We provide explicit formulas to compute the optimal cells, a complete characterization, and links between both resolution measures.

When applied to CNES's sampling systems, the resolution and information measures we propose provide a tool for analyzing the trade-offs posed by the different sampling systems. Moreover our optimal reciprocal cells may take very anisotropic shapes, and they consist sometimes of several connected components, far beyond the usual Voronoi (square or hexagonal) cell. In such cases, the only way to correctly represent all the information contained in the samples is to zoom (oversample) the image, and extrapolate its spectrum, in order to fill in the gaps between connected components in the spectral domain, thus avoiding ringing artifacts. To do so, we propose a modification of the data-fitting term of the deconvolution and zoom method proposed by Malgouyres and Guichard [122], which takes into account the shape of our optimal reciprocal cell. Our experiments show that this simple modification may either significantly reduce aliasing, or significantly improve resolution with respect to known deconvolution methods. If we want to restrict ourselves to linear deconvolution/denoising filters, our framework leads to a natural way to generalize the classical Wiener filter, so that it takes into account not only white noise but also aliasing.

Perturbed Sampling

Satellite images are not sampled on an exactly regular grid, but rather on a slightly perturbed grid. The sources of these perturbations include: micro-vibrations of the satellite while it takes the image, and irregularities in the position of the sensors on the image plane. For certain satellite images, the combined effect of these perturbations can be automatically estimated for each image, by different means developed at CNES. Therefore, in this chapter the perturbation will be assumed to be known, and we address the question of how to restore the regularly sampled image from the irreg-

ular samples, and under what conditions this restoration can be performed in a stable way. From this viewpoint we analyze the sensitivity of different sampling systems to perturbations in the sampling grid.

The chapter may be divided in two parts. In the first part we review the mathematical theory, which provides limits on the perturbation, beyond which a stable reconstruction is no longer possible. Satellite image sampling systems exhibit certain peculiarities, in particular the image is sampled at the critical (or sub-critical) sampling rate, but its spectrum near the Nyquist frequency may be relatively weak. Furthermore, the sampling grid is not a general irregular grid but a small perturbation of a regular grid, where the perturbation is very smooth (has a small spectral support) with respect to the image. Unfortunately none of the classical mathematical results on stable reconstruction from irregular samples takes all of these constraints into account ; they are excessively restrictive in one aspect or the other to be applicable to image sampling systems. Results based on Kadec's $\frac{1}{4}$ theorem do allow for critical sampling, but are too restrictive concerning the amplitude of perturbations. On the other hand, Beurling-Landau-type results take smoothness of the perturbation into account, but require a much too high sampling density.

In the second part we review the algorithms that are available for reconstruction from regular samples. As in the case of theoretical results, none of the available algorithms ensures its convergence to an exact solution under the conditions of satellite imagery, and they do not fully exploit its peculiarities. Nevertheless, two of them provide reasonable approximations, below the noise level at a quite high computational cost. For this reason, we propose a new algorithm, that better exploits the special characteristics of perturbed satellite imagery. Under favorable conditions it produces as good an approximation as the currently available methods, with about 5 times less computational effort. In several strongly perturbed instances, the accuracy of the new algorithm is, however, poorer than that of existing ones.

Bernard Rougé conjectured that an exact stable reconstruction of a bandlimited image is possible, when the image has weak high frequencies, and has been sampled on a perturbed grid, when the grid perturbation is much smoother than the image itself. This question remains open. Our experiments suggest that the answer to this question may be negative in the general case, and a more sensible question to ask is whether stable reconstruction *up to a certain accuracy* is possible under the same conditions, and how to determine this attainable accuracy level from the image and the perturbation.

Interpolation of Digital Terrain Models

Interpolation of Digital Elevation Models becomes necessary in many situations. For instance when constructing them from contour lines (available *e.g.* from non-digital cartography), or from disparity maps based on pairs of stereoscopic views, which often leaves large areas where point correspondences cannot be found reliably.

The Absolutely Minimizing Lipschitz Extension (AMLE) model is singled out as the simplest interpolation method satisfying a set of natural requirements. In particular, a maximum principle is proven, which guarantees not to introduce unnatural oscillations which is a major problem with many classical methods.

We then discuss the links between the AMLE and other methods, previously proposed in the geoscience literature.

In particular we show its relation with geodesic distance transforms, which are commonly used to interpolate level curves. We show that with a particular choice of the weights or potential function, iterated geodesic distance transforms quickly converge to AMLE under some rather restrictive assumptions. In the general case, however, we provide some evidence that the equivalence is not valid, even though possibly useful from the numerical viewpoint.

We also relate the AMLE to the thin plate method, that can be obtained by a prolongation of the axiomatic arguments leading to the AMLE, and addresses the major disadvantage of the AMLE model, namely its inability to interpolate slopes as it does for values. Nevertheless, in order to interpolate slopes, we have to give up the maximum principle and authorize the appearance of oscillations. We also discuss the possible link between the AMLE and Kriging methods that are the most widely used in the geoscience literature.

We end by numerical comparison between the different methods. Our experiments show that despite its inability to interpolate slopes, the AMLE produces qualitatively better results than other methods, since it better preserves ridges. Hence we propose a simple combination of the AMLE and Kriging models which preserves the slope-interpolating and ridge-preserving capabilities of Kriging and AMLE respectively. Either the combined method, or the AMLE, showed the best approximation results in all cases, both in terms of L^2 distances or qualitative properties.

Vanishing Point Detection

Vanishing points in digital images result from the projection by a pin-hole camera of a set of parallel lines in 3D. Most of the proposed computational methods for detection of vanishing points are forced to rely heavily on additional properties (like orthogonality or coplanarity and equal distance) of the underlying 3D lines or knowledge of the camera calibration parameters in order to avoid spurious responses. Furthermore, to the best of our knowledge, the question of determining whether an image contains vanishing points or not, and the exact number has not been systematically addressed. In this last part, we show that in many images of man-made environments vanishing points alone (without any additional assumption on the underlying 3D lines beyond pure parallelism) are meaningful Gestalts, in the sense of Helmholtz's perception principle, recently proposed for computer vision by Desolneux et al. [62], Desolneux, Moisan, and Morel [64]. This leads to a vanishing point detector with a low false alarms rate and a high precision level, which doesn't rely on any a priori information on the image or calibration parameters, and doesn't require any parameter tuning.

The key differences of this detector with respect to previous approaches are the following: *(i)* The line segments are themselves detected with almost-zero false alarms rate, by a refinement of the method proposed by Desolneux, Moisan, and Morel [61], also based on Helmholtz's principle; *(ii)* The threshold to determine a meaningful vanishing point from a large vote in the generalized Hough transform is computed in such a way that it guarantees a low false alarms rate; *(iii)* Finally a Minimum Description Length (MDL) criterion is used to further restrict the number of spurious vanishing points.

We also discuss two closely related generalizations of the MDL principle: one of them may serve to avoid the mutual exclusion of two vanishing points, the second one may be used to avoid, if desired, the masking of a "weak" vanishing point consisting of a small number of converging lines, by a "stronger" one where a huge number of lines converge.

Our experiments show that in urban scenes, the number and positions of vanishing points are correctly estimated. Other experiments on natural images, where accidental vanishing points are often detected, suggest (and further support the observation in [63], on) the need for modelling conflicts between different Gestalts, in this case between line segments and curved boundaries.

Notations

Set Algebra

$A \cup B$	union of two sets
$A \cap B$	intersection of two sets
A^c	complement of a set (the set containing all the elements not in A)
$\overset{\circ}{A}$	interior of set A
\overline{A}	closure of set A
∂A	frontier of set A
Range(f) or $R(f)$	the range of a function $f : A \rightarrow B$, <i>i.e.</i> the image $f(A)$ of the function's domain

Probability

$E\{X\}$	expected value of random variable X
$\text{Var}\{X\}$	variance of random variable X
$\text{Cov}\{X, Y\}$	covariance of random variables X and Y (a matrix if both are vectors)
$H(X)$	the entropy of a random variable X

Linear Algebra

A^T	matrix transpose
\overline{A}	complex conjugate of each element in matrix A
A^*	adjoint operator (transpose and complex conjugate in the case of matrices)
$\langle x, y \rangle$ or $x \cdot y$	scalar product of x and y

Mathematical Morphology

$A \oplus B$	dilation of A by structuring element B
$A \ominus B$	erosion of A by structuring element B
$A \bullet B$	closing of A by structuring element B
$A \circ B$	opening of A by structuring element B

Fourier Analysis

$\mathcal{F}(f) = \hat{f}$

$\overline{\mathcal{F}}(f)$

$\mathcal{F}^{-1}(f)$

$f * g$

h

$H = \hat{h}$

D^*

$|D^*|$

Γ

Γ^*

$\det(\Gamma)$

$\delta = \delta_0$

δ_γ

Δ_Γ

Differential calculus

$Df = \nabla f$

$D_i f$

$D^2 f$

$D^2 f(u, v)$

Fourier transform of f “inverse” Fourier transform of f , such that $\overline{\mathcal{F}}(\mathcal{F}(f)) = (2\pi)^d f$ inverse Fourier transform of f , such that $\mathcal{F}^{-1}(\mathcal{F}(f)) = f$ convolution of functions or distributions f and g

point-spread function

modulation transfer function

reciprocal cell

the area of the reciprocal cell

regular sampling grid

dual grid

refers actually to $\det(B)$ where B is a matrix whose columns are the set of basis vectors generating Γ . Observe that $\det(\Gamma) = |D|$ for any tile D of Γ .

Dirac delta

Dirac delta centered at point γ Dirac comb $\sum_{\gamma \in \Gamma} \delta_\gamma$ gradient of f derivative with respect to the i -th argument of f

second order derivative matrix

second order derivative with respect to vectors u and v , *i.e.* the matrix product $v^T \cdot (D^2 f) \cdot u$.

Part I

Image Sampling and Shannon Interpolation

Chapter 1

Measuring and Improving Resolution of Regularly Sampled Images; Applications to Restoration *

Abstract

Traditionally, discrete images are assumed to be sampled on a square grid and from a special kind of band-limited continuous image, namely one whose Fourier spectrum is contained within the rectangular “reciprocal cell” associated with the sampling grid. With such a simplistic model, resolution is just given by the distance between sample points.

Whereas this model matches to some extent the characteristics of traditional acquisition systems, it doesn’t explain aliasing problems, and it is no longer valid for certain modern ones, where the sensors may show a heavily anisotropic transfer function, and may be located on a non-square (in most cases hexagonal) grid.

In this chapter we first summarize the generalizations of Fourier theory and of Shannon’s sampling theorem, that are needed for such acquisition devices. Then we explore its consequences: (i) A new way of measuring the effective resolution of an image acquisition system; (ii) A more accurate way of restoring the original image which is represented by the samples. We show on a series of synthetic and real images, how the proposed methods make a better use of the information present in the samples, since they may drastically reduce the amount of aliasing with respect to traditional methods. Finally we show how in combination with Total Variation minimization, the proposed methods can be used to extrapolate the Fourier spectrum in a reasonable manner, visually increasing image resolution.

*This chapter was submitted for journal publication jointly with Sylvain Durand and Bernard Rougé, in October 2002. A preliminary version was published as a CMLA preprint [17].

1 Introduction

When an analog image is converted into digital form it is sampled on a regular grid. If the image is not band-limited at the right frequency, this sampling operation produces a replication of higher frequency components into lower frequency components of the spectrum, a phenomenon known as “aliasing”. Aliasing artifacts in digital images and sampling grids other than the traditional square grid have received little attention in the mainstream literature on image processing.

Some textbooks on image processing like [40, 145] do introduce the formalisms needed for sampling in general lattices, they describe the corresponding aliasing problems and the corresponding generalization of Shannon’s sampling theorem (we present it here in a similar fashion). Most commonly, however, sampling and aliasing are only described in the case of rectangular grids and reciprocal cells [84, 100, 153] or at most in the case of quincunx grids and interlaced television signals [101]. In any case in all these textbooks the treatment of non-orthogonal sampling grids and aliasing is limited to a description of the effects and the tools to analyze them, but rarely gives any practical means of measuring or reducing the amount of aliasing for a given sampling system, maybe with the exception of [139] which gives different measures of resolution, aliasing and interpolation error, as well as a Wiener filtering technique which takes into account the aliasing noise. Otherwise, the degradation model assumed in the great majority of the image restoration literature, includes convolution with a band-limited kernel plus noise, and thus ignores the effects of aliasing. Discussions on how to limit the aliasing error are far more common in the computer graphics literature, where one has full control of the sampling system (most notably to its sampling rate and transfer function), thus leading to approaches that are in most cases not applicable in digital imaging, where both parameters are dictated by physical constraints of the imaging optics and the electronics of the sensors.

Concerning sampling on non-orthogonal grids, most textbooks dedicate a section to compare the properties of hexagonal and square grids, but then present the rest of the results only for the square grids. More particularly, concerning Fourier analysis in hexagonal grids, an important part of the literature started by [128] and followed by many others [68, 129, 176], deals with specialized FFT algorithms for the hexagonal sampling grids. Whereas this is necessary if we want to keep a rectangular periodization pattern, if we consider the image to be periodic along the same axes as the sampling grid, then we can still use the usual FFT [135] as we do in square grids. We

believe that the little emphasis that has been given to this property as well as the lack of specialized literature discouraged the development of hexagonal image processing despite its advantages.

In fact, most authors agree on several advantages of the hexagonal grid over the more traditional rectangular grid. We summarize the most important ones here :

- *Sampling density*: under isotropic conditions a hexagonal grid requires 13.4% less samples than a square grid to represent the same amount of information.
- *Discrete topology*: we only need a single natural connectivity concept (6 neighbours), which is consistent with Jordan's curve theorem, whereas in the square grid consistency with the continuous case obliges us to switch between 4-connectivity for the foreground and 8-connectivity for the background or conversely. This greatly simplifies many morphological operations.
- *Isotropic image analysis*: It has been shown that in the hexagonal grid we can construct more isotropic wavelet families with much better orientation selectivity [49, 56, 97, 112, 114, 130, 157, 161, 162] , more accurate distance transforms [38] , isotropic 3-separable filters, etc.

For a more thorough discussion on the properties of hexagonal grids see chapter A and [5, 168] .

Coming back to aliasing, despite the little importance given to the subject in the image processing literature, it has been recognized by several authors as an important problem, which significantly limits the quality of image data [34] . More recently, the importance of aliasing artifacts present in most DVD movies was analyzed [133] , and the physical constraints of any CCD array leading to a significant amount of aliasing were described in [147] . At the same time the latter work as well as [116, 131, 149] propose a physical sampling device that allows to avoid the limitations of CCD arrays and reduce aliasing. Such sampling devices involve any kind of regular sampling grid (rectangular, hexagonal, or in general, any grid generated by two independent vectors on the plane), and may show heavily anisotropic transfer functions.

The main purpose of this chapter is to stress once more the importance of aliasing in imaging systems, provide ways of measuring to what an extent it may degrade the image quality, and provide practical methods to reduce the amount of aliasing of any sampling system, provided we know its three main characteristics: the geometry of the sampling grid, the transfer function, and the noise. The results can be used both by the designer of imaging systems who wants to choose between different possibilities the

one which provides the best image-quality/cost ratio, as well as by the image processing engineer who wants to obtain the best possible reconstructed image from a given system.

We start by giving in section 2.2 a slightly generalized form of Shannon's sampling theorem, since:

1. It applies to any regular sampling grid, not necessarily rectangular.
2. Instead of the classical rectangular (Voronoi) reciprocal cell it allows more general shapes of this cell which is used both to express the band-limiting condition and the sinc-interpolation formula.

This form has long been known by the mathematicians and similar forms can be found in [36, 80, 135, 145], but its flexibility has not been fully exploited by the image processing community yet.

Next, we review in sections 2.3 and 2.4 the common notions of resolution, critical sampling and sampling efficiency, and extend them to the new formulation. We also present some simple examples on how to adapt the reciprocal cell to the image contents in the ideal case of strictly band-limited images and no noise.

In section 3 we describe the main characteristics of some of the sampling systems mentioned above: both the classical CCD arrays, and the more anisotropic systems described in [116, 131, 149]. Then in sections 4 and 5 we extend to this more realistic case the ideas of resolution and optimal reciprocal cell that were introduced in sections 2.3 and 2.4 in the ideal band-limited case.

Finally in section 6 we present the results of our numerical experiments, commenting the quantitative comparison of the different systems according to the resolution-based measures that we introduced, and we propose an application of the optimal reciprocal cell to image restoration and spectral extrapolation, which considers not only convolution and noise, but also aliasing within the degradation model.

Before we start with the main presentation let us present a simple example that illustrates how important it may be to consider aliasing during restoration. We want to sample a square wave signal f with period $A = 12$, with a sampling system consisting of an aperture of $T \approx 2.38$, and a sampling rate of 1, which adds a white noise with standard deviation 0.01 (*i.e.* 1% of the wave's amplitude). Summarizing the sampled signal will be

$$g(k) = (f * h)(k) + n(k) \text{ for all } k \in \mathbb{Z} \quad (1.1)$$

where the point-spread function is $h(t) = \mathbb{1}_{[-\frac{T}{2}, \frac{T}{2}]}(t)$. Figure 1.1 shows the different





(a) original signal f and its Fourier transform(b) pre-sampling filter h and transfer function \hat{h} (c) sampled signal $g = f * h + n$ and sinc-interpolation(d) Wiener-filter restoration f_1 with voronoi reciprocal cell(e) Wiener-filter restoration f_2 with optimal reciprocal cell

Figure 1.1: Example of Wiener filter restoration with aliasing. In (d) the spectrum of the restored image is $\hat{f}_1(\xi) = \hat{g}(\xi)K(\xi)$ for $|\xi|$ smaller than the Nyquist frequency and zero otherwise. Here $K = \frac{H^*}{|H|^2 + \sigma^2}$ is the Wiener deconvolution filter and $H = \hat{h}$ is the modulation transfer function (MTF). Observe that this produces a bad result because we are deconvolving an aliased coefficient near the zero-crossing of H . In (e) we solve this problem by deconvolving $\hat{f}_2(\xi) = \hat{g}(\xi)K(\xi)$ for ξ within an optimal reciprocal cell adapted to the MTF, and $\hat{f}_2(\xi) = 0$ for ξ outside this reciprocal cell.

stages in the sampling of g and its restoration with a Wiener filter. Observe that the result is catastrophic for the following reason, that is more clearly explained in figure 1.2: The MTF (the modulation transfer function *i.e.* the Fourier transform $H = \hat{h}$ of the blurring kernel h) has a zero crossing at a frequency $\xi_0 \approx 0.4$ just before the Nyquist frequency. If there was no aliasing, the corresponding Fourier coefficient of the sampled image $\hat{g}(\xi_0)$ should be nearly zero, plus noise. But since the MTF has a second lobe beyond the Nyquist frequency, some high frequency coefficients (see for instance the fifth peak at frequency $\xi_0 - 1 \approx -0.6$) are aliased near frequency ξ_0 ¹. Thus $\hat{g}(\xi_0)$ is far larger than the expected signal plus noise level, and when we try to deconvolve it by $\frac{H^*}{|H|^2 + \sigma^2}(\xi_0) \approx \frac{0.01}{0.01^2 + 0.01^2} = 50$ it blows up.




Figure 1.2: Example of Wiener filter restoration with aliasing (continued). Left: Closeup of the system's transfer function \hat{h} , the spectrum \hat{f} of the original signal, the voronoi reciprocal cell (consisting of all frequencies smaller than the Nyquist frequency associated to the sampling rate), and the optimal reciprocal cell, adapted to the MTF. To the right we show the (periodic) spectrum of the sampled signal with the two Wiener deconvolution filters. Observe how the optimal reciprocal cell (and the corresponding deconvolution filter) avoids a region around the zero-crossing of \hat{h} at $\xi_0 \approx 0.4$ and substitutes it by the region around its alias at $\xi_0 - 1$, where the magnitude of \hat{h} is much larger.

This suggests that we should avoid deconvolving around ξ_0 , and do it rather around $\xi_0 - 1$ where the real information comes from. Put another way, we defined an alternative reciprocal cell, which is still compatible with Shannon's sampling theorem, but is better adapted to our instrument's transfer function. In the following sections we shall

¹ The magnitude of the non-aliased part of $\hat{g}(\xi_0)$ is $|\hat{f}(\xi_0)\hat{h}(\xi_0)| + \sigma = 0.84 * 0.01 + 0.01 = 0.018$, but the actual aliased coefficient is $\hat{g}(\xi_0) = 0.13$, seven times larger. The largest contribution to $\hat{g}(\xi_0)$ comes from frequency $\xi_0 - 1 \approx -0.6$, *i.e.* $|\hat{f}(\xi_0 - 1)\hat{h}(\xi_0 - 1)| = 0.58 * 0.20 = 0.12$

develop a general technique for choosing this reciprocal cell in an optimal way.

With the adapted Wiener deconvolution we obtain a reconstruction error $\frac{\|f_1 - f\|}{\|f\|}$ of 20%, which is only slightly better than the original 24% error in the sampled image g , but it is much better than the 146% reconstruction error of the original Wiener filter. The ringing artifacts in the result are mainly due to the Fourier coefficients that were lost between the two lobes of the optimal reciprocal cell. To avoid it we need to extrapolate the spectrum between these two lobes. In section 6 we propose to do so by minimizing the total variation.

2 Review of Fourier Analysis in non-square grids

In this section we recall the mathematical basics of sampling theory in the ideal case of infinite spatial support and compact spectral support (band limited). In the following section 3 we will discuss how this modeling applies to real acquisition systems. The main elements of this review are contained in [36, 80, 135], but here we give slightly different formulations which are better suited for defining our image resolution measures.

Section 2.1 defines the geometrical bases of sampling in non-orthogonal grids and section 2.2 provides a formal proof of Shannon's sampling theorem on such grids. A more rigorous proof, as well as different formulations depending on the degree of regularity of the image are given in appendix A.

Finally sections 2.3 and 2.4 review some common concepts related to image resolution in the ideal case of compact spectral support and no noise. These concepts will be generalized in sections 4 and 5 to the more realistic systems described in section 3.

2.1 Sampling grids and tilings of the plane

We will consider our discrete images to be defined on a *regular sampling grid*

$$\Gamma := \{n_1 e_1 + n_2 e_2 : n_1, n_2 \in \mathbb{Z}\} = \mathbb{Z}e_1 + \mathbb{Z}e_2 \quad (1.2)$$

where $\{e_1, e_2\}$ is a basis of \mathbb{R}^2 . To each sampling grid we shall associate a *dual grid*

$$\Gamma^* := \{n_1 e_1^* + n_2 e_2^* : n_1, n_2 \in \mathbb{Z}\} = \mathbb{Z}e_1^* + \mathbb{Z}e_2^* \quad (1.3)$$

where $\langle e_i^*, e_j \rangle = 2\pi\delta_{ij}$

This will be useful when we define the Fourier transform of a discrete image.

In the sequel we do not impose any further condition on the generating vectors e_1 and e_2 . In practice, however, only two such grids are used, namely:

$$\Gamma_4 = \mathbb{Z} \begin{pmatrix} 1 \\ 0 \end{pmatrix} + \mathbb{Z} \begin{pmatrix} 0 \\ 1 \end{pmatrix} \quad (\text{square grid})$$

$$\Gamma_6 = \mathbb{Z} \begin{pmatrix} 1 \\ 0 \end{pmatrix} + \mathbb{Z} \begin{pmatrix} 1/2 \\ \sqrt{3}/2 \end{pmatrix} \quad (\text{hexagonal grid}).$$

The reason is that these are the only ones which generate tilings of the plane by regular polygons, and show the highest number of symmetries. The corresponding dual grids are

$$\Gamma_4^* = 2\pi \left(\mathbb{Z} \begin{pmatrix} 1 \\ 0 \end{pmatrix} + \mathbb{Z} \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right) \quad (\text{dual square grid})$$

$$\Gamma_6^* = 2\pi \frac{2}{\sqrt{3}} \left(\mathbb{Z} \begin{pmatrix} \sqrt{3}/2 \\ -1/2 \end{pmatrix} + \mathbb{Z} \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right) \quad (\text{dual hexagonal grid})$$

as shown in figure 1.3.

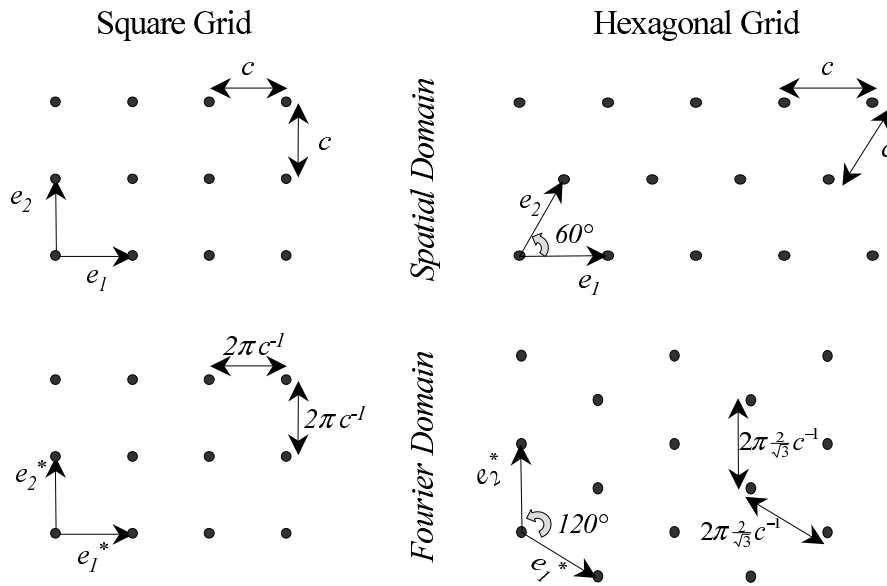


Figure 1.3: The most common sampling grids and their corresponding duals. The square and hexagonal grids can be singled out as the only two possibilities from a series of invariance properties [5]. As explained in section 2.2, when an image is sampled on a grid Γ , its Fourier transform is periodic with respect to its dual grid Γ^* , whose generating vectors are biorthogonal to those of Γ .

Usually, sensor arrays are composed of a set of sensors all with the same shape and characteristics, with their centers lying on a regular grid, and covering the image plane

without superposition. More precisely, given a set $D \subset \mathbb{R}^2$ and a grid Γ we shall say that:

- (Γ, D) is a *covering*², if $\bigcup_{\gamma \in \Gamma} (D + \gamma) = \mathbb{R}^2$
- (Γ, D) is a *packing*², if $\forall \gamma \in \Gamma, \gamma \neq 0 \implies D \cap (D + \gamma) = \emptyset$
- (Γ, D) is a *tiling*, if it is both a packing and a covering.

Here $D + \gamma$ represents the translated set $\{x + \gamma : x \in D\}$. We slightly relaxed the usual covering and packing conditions by considering the closure \bar{D} and the interior $\overset{\circ}{D}$ of the set D . We do so for clarity, to simplify the definitions of our sets D (otherwise, we need a convention to decide whether to include or not in D a point in the border ∂D). We shall also say that D is a *tile* of the grid Γ if (Γ, D) is a tiling. Similarly, we say that D is a *sub-tile* (resp. a *super-tile*) if (Γ, D) is a packing (resp. a covering). These three concepts, when applied to the dual grid Γ^* will also be useful to choose suitable spectral domains.

In figure A.3 we showed the only two tilings of the plane where the tile is a regular polygon. The triangular tile can only cover the plane if –in addition to the integer translations in Γ – we also allow an horizontal symmetry. We also observed that it has the smallest symmetry group, whereas the hexagonal tiling possesses the largest one. For this and other reasons the triangular grid is not very useful in image processing. In the sequel we shall mostly concentrate on the square and hexagonal sampling grids, but will occasionally use tiles for these grids, which are not regular polygons. Recall that in figure A.2 we showed an example, which illustrates how one can construct such tilings, without changing the sampling grid Γ .

2.2 Shannon's sampling theorem.

To each sampling grid we shall associate a *Dirac comb*

$$\Delta_\Gamma := \sum_{\gamma \in \Gamma} \delta_\gamma \quad (1.4)$$

composed of delta-functions, such that $\langle \delta_\gamma, f \rangle = \int f(x) \delta_\gamma(dx) = f(\gamma)$. Sampling on a grid Γ can then be expressed as simply multiplying by Δ_Γ .

A sampled image will be modeled as a Γ -discrete distribution, i.e. a tempered distribution g for which there exists a continuous function $f \in \mathcal{O}_M(\mathbb{R}^2)$ such that $g = \Delta_\Gamma \cdot f$.

² Here, set equality is to be understood in the weak sense or “almost everywhere”, namely $A = B$ iff both sets $A \setminus B$ and $B \setminus A$ have zero measure.

Similarly, we will say that a tempered distribution g is Γ -periodic if it can be expressed as a convolution $g = \Delta_\Gamma * f$ for some tempered distribution with fast decay $f \in \mathcal{O}'_C(\mathbb{R}^2)$.

Observe that this definition is equivalent to the usual definition of periodicity for functions, namely that $g(x + \gamma) = g(x)$ for all $x \in \mathbb{R}^2, \gamma \in \Gamma$.

In fact if $g = \Delta_\Gamma * f$, then $g(x + \gamma_0) = \sum_{\gamma \in \Gamma} f(x + \gamma + \gamma_0) = \sum_{\gamma \in \Gamma} f(x + \gamma)$, since the mapping $\gamma \mapsto (\gamma + \gamma_0)$ is bijective in Γ . Conversely if $g(x + \gamma) = g(x)$ then $g = \Delta_\Gamma * (\chi \cdot g)$ where $\chi \in C_0^\infty(\mathbb{R}^2)$ is a partition of unity (i.e. such that $\sum_{\gamma \in \Gamma} \chi(x + \gamma) \equiv 1$).

It is easy to show that such a partition of unity exists. It suffices to start from a non-negative function $\varphi \in C_0^\infty$ which is strictly positive on a tile D of Γ , and then define $\chi(x) = \frac{\varphi(x)}{\sum_{\gamma \in \Gamma} \varphi(x + \gamma)}$.

We recall the continuous Fourier Transform of $f \in L^1$ and its inverse

$$\hat{f}(\xi) = \mathcal{F}(f)(\xi) := \int_{\mathbb{R}^2} e^{-i\langle x, \xi \rangle} f(x) dx \quad (1.5)$$

$$(2\pi)^2 f(x) = \overline{\mathcal{F}}(\hat{f})(x) := \int_{\mathbb{R}^2} e^{i\langle x, \xi \rangle} \hat{f}(\xi) d\xi \quad (1.6)$$

as well as some of its properties that we state without proof:

- $\overline{\mathcal{F}}(\mathcal{F}(f)) = (2\pi)^2 f$
- $\mathcal{F}(f * g) = \mathcal{F}(f) \cdot \mathcal{F}(g)$
- $\mathcal{F}(f \cdot g) = \frac{1}{(2\pi)^2} \mathcal{F}(f) * \mathcal{F}(g)$
- $\mathcal{F}(\Delta_\Gamma) = |D^*| \Delta_{\Gamma^*}$ where D^* is a tile of Γ^* .

These properties are not valid in general. The precise conditions, as well as some extensions of the Fourier transform and its properties to more general classes of functions and distributions is given in appendix A.1. Here we use these properties to give a formal derivation of the sampling theorem. For a rigorous proof and some variants of its formulation please refer to appendix A.2.

From the last two properties we deduce that the Fourier transform of a Γ -discrete image g is Γ^* -periodic:

$$\hat{g} = \mathcal{F}(\Delta_\Gamma \cdot f) = \frac{|D^*|}{(2\pi)^2} \Delta_{\Gamma^*} * \hat{f}. \quad (1.7)$$

Thus \hat{g} is completely determined if we know its value on a tile D^* of the dual grid Γ^* . (In fact, since D^* is a tile, for any $\xi \in \mathbb{R}^2$ we can find $\gamma^* \in \Gamma^*$ such that $(\xi + \gamma^*) \in D^*$, and therefore by periodicity $g(\xi) = g(\xi + \gamma^*) = g|_{D^*}(\xi + \gamma^*)$). If in addition $\hat{g}|_{D^*} = \hat{f}|_{D^*}$ then g is completely determined by $\hat{f}|_{D^*}$.

A slightly stronger condition than $\hat{g}|_{D^*} = \hat{f}|_{D^*}$ ensures the converse, *i.e.* that the continuous image f is completely determined by the discrete image $g = \Delta_\Gamma \cdot f$. It suffices to require that f be *band-limited*, *i.e.* that $\text{supp}(\hat{f}) \subseteq D^*$ with D^* a sub-tile of Γ^* . In that case $\mathbb{1}_{D^*} \cdot (\Delta_{\Gamma^*} * \hat{f}) = \mathbb{1}_{D^*} \cdot \hat{f} = \hat{f}$.³ This allows us to recover \hat{f} from \hat{g} as follows:

$$\frac{(2\pi)^2}{|D^*|} \mathbb{1}_{D^*} \cdot \hat{g} = \frac{(2\pi)^2}{|D^*|} \mathbb{1}_{D^*} \cdot \left(\frac{|D^*|}{(2\pi)^2} \Delta_{\Gamma^*} * \hat{f} \right) = \mathbb{1}_{D^*} \cdot (\Delta_{\Gamma^*} * \hat{f}) = \hat{f}. \quad (1.8)$$

Finally we apply the Fourier inversion formula to obtain:

$$\frac{1}{|D^*|} \overline{\mathcal{F}}(\mathbb{1}_{D^*}) * g = f \quad (1.9)$$

This formally shows a generalized form of Shannon's sampling theorem:

Theorem 1 (Shannon-Whittaker). *Consider a function $f \in L^2(\mathbb{R}^d)$, a d -dimensional sampling grid Γ and a compact reciprocal cell $D^* \subset \mathbb{R}^d$. If the following conditions are met*

(S1) *the reciprocal cell D^* is a sub-tile of the dual grid Γ^**

(S2) *$\text{supp}(\hat{f}) \subseteq D^*$ (i.e. f is band-limited),*

then f can be completely recovered from its samples in Γ , i.e. from

$$g = \Delta_\Gamma \cdot f. \quad (1.10)$$

In fact, $g \in l^2(\Gamma)$, and the reconstruction is given by the following convolution

$$f(x) = \sum_{\gamma \in \Gamma} g_\gamma s(x - \gamma) \quad (1.11)$$

of the sampled image g with the generalized sinc kernel $s = \frac{1}{S^} \overline{\mathcal{F}}(\mathbb{1}_{D^*})$. This equality holds in general only in the sense of the L^2 norm. If in addition $g \in l^1(\Gamma)$, then there is pointwise equality and uniform convergence.*

2.3 Choice of the reciprocal cell.

The main difference between this formulation of Shannon's sampling theorem and the more classical one is (apart from the fact that we are using a possibly non-orthogonal sampling grid) the introduction of one additional degree of freedom, namely the reciprocal cell D^* , which has to meet both conditions (S1) and (S2). As long as both conditions are met the reconstruction formula is valid, but we are free to choose any of them. In practice D^* is most commonly assumed to be in one of the two cases:

³ Since D^* is a sub-tile of Γ^* , if $\xi \in D^*$ then $\xi + \gamma^* \notin D^*$ for $\gamma^* \in \Gamma^* \setminus \{0\}$. Now, since $\text{supp}(f) \subseteq D^*$, only the term in $\gamma^* = 0$ is non-zero in $\sum_{\gamma^* \in \Gamma^*} \hat{f}(\xi + \gamma^*)$. This shows that $(\Delta_{\Gamma^*} * \hat{f})|_{D^*} = \hat{f}|_{D^*}$. Consequently $\mathbb{1}_{D^*} \cdot (\Delta_{\Gamma^*} * \hat{f}) = \mathbb{1}_{D^*} \cdot \hat{f}$. But again, since $\text{supp}(\hat{f}) \subseteq D^*$ we have $\mathbb{1}_{D^*} \cdot \hat{f} = \hat{f}$.

Spectral support of the image. That means we take $D^* = \text{supp}(\hat{f})$. This choice has the obvious advantage that it always satisfies (S2), but it presents several disadvantages. First there is no reason why this definition should lead to a convenient filter in the reconstruction formula (1.11). The alternative we discuss below addresses exactly this issue. Secondly, \hat{f} will likely have no compact support, so that actually we will be forced to consider an “essential support” where the energy of \hat{f} is large enough. We shall come back to this idea in the next section.

Voronoi reciprocal cell. The other canonical choice consists of choosing D^* to be the largest and most isotropic cell that satisfies (S1), namely we choose as a reciprocal cell the Voronoi set $D_{\text{Vor}}^* = \text{Vor}(\Gamma^*)$ of the dual grid, as defined below.

$$D_{\text{Vor}}^* = \text{Vor}(\Gamma^*) = \{\gamma^* \in \mathbb{R}^d : \forall \gamma_0^* \in \Gamma^*, \gamma_0^* \neq 0 \Rightarrow \|\gamma^*\| < \|\gamma^* + \gamma_0^*\|\} \quad (1.12)$$

It can be easily verified that D_{Vor}^* is an admissible tile for the dual grid Γ^* , hence it always satisfies condition (S1) in Shannon’s sampling theorem. In addition its mathematical and numerical treatment is much easier. In fact, in dimension 2, the Voronoi cell can be expressed as the product of three one-dimensional indicator functions, thus the spatial-domain convolution is “three-separable” in the sense that it can be expressed as a sequence of three one-dimensional convolutions. (In the case of the square grid, the third convolution would be in the direction of the diagonal and becomes unnecessary). Figure 1.4 shows the Voronoi cell corresponding to the square and hexagonal grids, plus an intermediate case where the basis vectors of the grid form an angle of 80° .

Whereas the Voronoi reciprocal cell provides more convenient formulas and always satisfies (S1), it does not address the band-limiting condition (S2). In order to do so, it should be adapted to the spectral support of the image $\text{supp}(\hat{f})$. In fact as we will see in section 5, D_{Vor}^* is the most sensible choice whenever \hat{f} is expected to be isotropic and radially decreasing, which is often the case in many imaging systems. However, many of the high resolution systems currently under development and study at CNES and other research centers, as well as some other more common situations do not satisfy these conditions (see section 3), which motivated to a large extent the present study.

But before going into the more complicated resolution measures, let us start with a few simple examples of systems that would prefer a Voronoi reciprocal cell, and others that would require another kind of D^* . These simple examples will be enough to show some of the central ideas in the rest of the chapter. To do so we introduce a simple quantitative measure of sampling density and sampling efficiency, which is used to

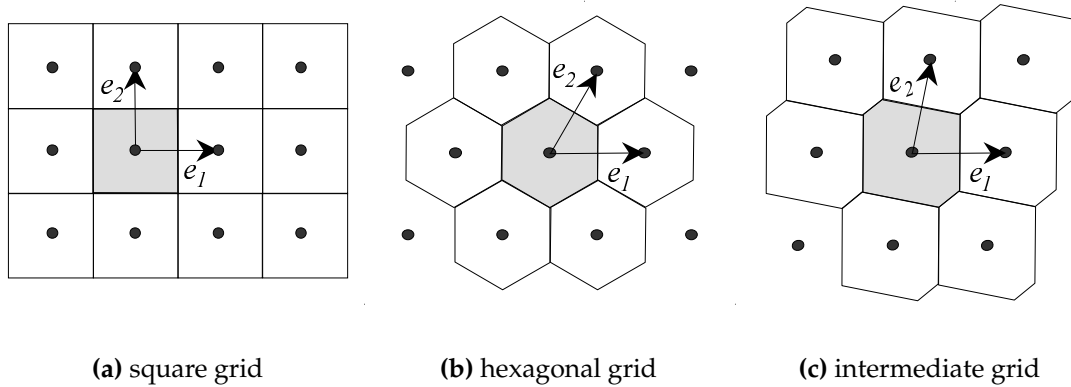


Figure 1.4: Voronoi reciprocal cell for three grids: square (generating vectors at 90°), hexagonal (60°), and an intermediate case where the generating vectors form an angle of 80° . Observe that the shape of this cell is in general a (not necessarily regular) hexagon, except in the square case. Hence the corresponding convolution filters in Shannon's reconstruction formula are in general 3-separable, whereas in the square case they are 2-separable.

give a precise meaning to phrases like: "this is the best sampling grid", or "this is the best reciprocal cell for this image".

2.4 Critical Sampling.

Definition 1 (Sampling density). The sampling density of a grid Γ is the number of samples per unit area, given by $\rho(\Gamma) = \frac{1}{\det(\Gamma)}$.⁴ Equivalently the sampling density can be computed as $\rho(\Gamma) = \frac{1}{|D|}$ for any admissible tile D .

One would like to satisfy Shannon's sampling theorem with as least samples as possible, *i.e.* with the lowest possible value of $\rho(\Gamma)$. The following definition expresses this desired limiting case:

Definition 2 (Critical sampling). Γ is a critical sampling grid for an image f iff

- (i) $(\Gamma^*, \text{supp}(\hat{f}))$ is a packing, and
- (ii) $\rho(\Gamma) = \inf \left\{ \rho(\Gamma_1) : (\Gamma_1^*, \text{supp}(\hat{f})) \text{ is a packing} \right\}$.

In such a case, since $(\Gamma^*, \text{supp}(\hat{f}))$ is a packing, there exists a reciprocal cell D^* satisfying both conditions in Shannon's theorem. The trade-off between sampling density

⁴ $\det(\Gamma)$ refers here to $\det(B)$ where B is a matrix whose columns are the set of basis vectors generating Γ . Observe that $\det(\Gamma) = |D|$ for any tile D of Γ .

and the packing condition becomes clearer if we express the sampling density directly in terms of the dual grid. In fact from the definition of dual grid we observe that

$$\rho(\Gamma) = \frac{1}{2\pi} \det \Gamma^* = \frac{1}{(2\pi)^2 \rho(\Gamma^*)} = \frac{|D^*|}{(2\pi)^2} \quad (1.13)$$

for any admissible tile D^* of the dual grid Γ^* . So we want a dual grid with the smallest possible tile, which still contains $\text{supp}(\hat{f})$. The ideal situation is when $\text{supp}(\hat{f})$ is itself a tile, which leads us to the following definition.

Definition 3 (Sampling efficiency). *The sampling efficiency of a grid Γ with respect to an image f , which satisfy Shannon's theorem is defined as $\frac{|\text{supp}(\hat{f})|}{(2\pi)^2 \rho(\Gamma)}$ or equivalently as*

$$\frac{|\text{supp}(\hat{f})|}{|D^*|} \quad (1.14)$$

for any tile D^* of the dual grid Γ^* .

Thus, for a given spectral support, the sampling efficiency is inversely proportional to the sampling density, and we obtain critical sampling when the sampling efficiency is maximal among all grids satisfying Shannon's theorem. The sampling efficiency actually coincides with the *packing density*, which has been studied in a more general setting, and in more dimensions. In dimension two the main result states that the maximal packing density with a circular tile is attained for the hexagonal grid (even if we allow much more general transformations that just integer multiples of two generating vectors) [27, 143, 174]. Restating this result in sampling terms we can say that when the spectral support of f is a circular domain, then the maximal sampling efficiency (critical sampling) is obtained when the dual sampling grid Γ^* is hexagonal. Figure A.6 and table A.1 illustrate how this happens and to what an extent the hexagonal sampling is more efficient than the square sampling in this case. From these figures we can conclude that a circular spectral support satisfying Shannon's theorem requires 13.4% more samples in the case of a square grid than in the case of an hexagonal grid.

Observe that despite the fact that the square grid is not the most efficient sampling in this case, the square dual sampling grid shown in figure A.6 is still the best possible choice among square grids. In this case we say that this square grid is critical within its geometry, but not globally critical:

Definition 4. Γ is a critical sampling grid for an image f within its geometry iff

- (i) $(\Gamma^*, \text{supp}(\hat{f}))$ is a packing, and

$$(ii) \quad \rho(\Gamma) = \inf \left\{ \rho(\Gamma_1) : (\Gamma_1^*, \text{supp}(\hat{f})) \text{ is a packing, and } \exists a \in \mathbb{R}, \Gamma_1 = a\Gamma \right\}.$$

In the examples discussed so far, in particular for both examples in figure A.6, the Voronoi reciprocal cell is admissible in Shannon's theorem, even in the case of critical or partially critical sampling. And this, due to the isotropy of the spectral support.

This is no longer the case in the following example, where the spectral support of f is assumed to be elliptic with eccentricity 2 and with the major axis at 45° from the horizontal. Assume that we want to sample an image with such a support with a vertical/horizontal rectangular sampling grid. If we stick to Voronoi reciprocal cells, then our most efficient sampling will be far from critical (even within geometry), as shown in figure 1.5(a). We can obtain a more efficient sampling by relaxing the sampling rate in the horizontal direction, thus leading to the reciprocal cell in figure 1.5(b) with the shape of a parallelogram. This is still not critical, since we can still relax a little bit the vertical sampling rate (though with a more complicated D^*), but this is enough for illustration purposes.

Observe that a correct choice of D^* is necessary only when applying Shannon's theorem to visualize, restore or interpolate the image. In the coding phase, any other reciprocal cell which forms a tiling would do just as well. The rectangular cell shown in figure 1.5(c) contains all the information about the original image. Some parts of the original spectral support were lost but they reappear as aliases in other regions, not used by the rest of the spectral support. If we display the sampled image as is (without any interpolation), we would in fact see these aliasing artifacts, but since we know the shape of the spectral support, we can relocate each coefficient to its correct position. To do so we need some spectral space of course, and we have to oversample at the rate given by the Voronoi cell in figure 1.5(a). This is one of the main ideas that will be exploited in sections 4 through 6, but in a more realistic case, with noise, with unbounded spectral support but known decay rate, and some other more realistic details that will be the subject of the next section.

3 Modeling of image acquisition systems

In the previous section we assumed an ideal situation, but real sampling systems present two limitations which force us to deal with approximations to Shannon's sampling theorem:

- *noise*, due to the fact that sampling systems can only make approximate measurements, and

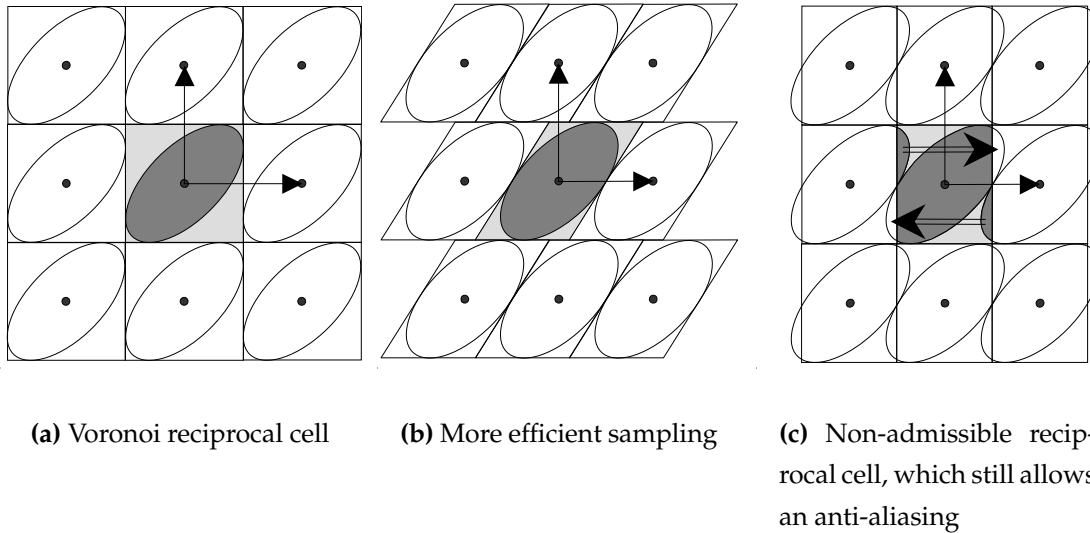


Figure 1.5: The figures above show the dual grid Γ and its generating vectors as well as the spectral support of the image to be sampled in dark-gray, and the chosen reciprocal cell in light-gray. We also show some aliases of the spectral support $\text{supp}(\hat{h}) + \gamma^*$ and the chosen reciprocal cell $D^* + \gamma^*$ for some $\gamma^* \in \Gamma^*$. Observe that Voronoi reciprocal cells (as in subfigure a) no longer yield a critical sampling when the spectral support is strongly anisotropic. To obtain a more efficient sampling we have to use a better adapted reciprocal cell like the parallelogram in subfigure (b). If displayed as-is (*i.e.* with no interpolation or with the usual sinc or zero-padding interpolation) the sampled image will show some aliasing artifacts as illustrated in subfigure c, but these can be solved by using an admissible reciprocal cell in the reconstruction formula. Equivalently, one can translate the aliased Fourier coefficients to their original location. The value of $\gamma^* \in \Gamma^*$ to be used for this translation can be deduced from the knowledge of the admissible reciprocal cell. Once the Fourier coefficients were translated to their correct location, we can resample at the rate given by the Voronoi reciprocal cell in subfigure (a).

- the fact that the analog image before sampling actually has *no compact spectral support*, and is neither *periodic* nor has an *infinite spatial support*.

In fact, a generic image sampling system can be viewed as a combination of:

- an optical system whose purpose is to focus the incoming light in the focal plane, and
- a set of sensors distributed on the focal plane which count the number of photons which meet this precise location during a certain period of time.

and the whole system might be moving during capture. The whole action of the sampling system can be modeled by the following formula:

$$g = \mathbb{1}_\Omega \cdot (\Delta_\Gamma \cdot (\overline{\mathcal{F}}(H) * f) + n) \quad (1.15)$$

where

- f represents the ideal image just before it enters the imaging system, this is assumed to be a radon measure;
- n represents the noise due to the imprecision in the sensor measurements;
- Γ represents the geometry of the array of sensors, which are assumed to be distributed on a regular grid;
- Ω represents the extent of the sensor array, which can never be infinite;
- H represents the transfer function of the instrument, which is a combination of several effects to be detailed later.

We observe that even in the absence of noise, and even if the transfer function H had a compact support, the resulting sampled image would have no compact spectral support, due to multiplication by $\mathbb{1}_\Omega$. Hence the periodization of the finite-extent sampled image g (needed to obtain a discrete spectrum) will introduce some aliasing. However this difficulty can be solved by a standard technique which consists of multiplying the whole image by a smooth window having a very smooth transition between 0 and 1 near $\partial\Omega$. With this technique this necessary amount of aliasing can be usually reduced to a level much below the noise level, so in the sequel we will assume a simplified image model:

$$g = \Delta_\Gamma \cdot (\overline{\mathcal{F}}(H) * f) + n \quad (1.16)$$

in which the sampled image g can be considered infinite or periodic.

Now let's consider the transfer function (abbreviated MTF and denoted by H). It can be modeled as the product of three main effects:

$$H = H_{\text{sen}}H_{\text{mov}}H_{\text{opt}}. \quad (1.17)$$

The first term is due to the fact that the sensors do not pick the value of f at one point but rather integrate all the photons that arrive at its sensitive region which usually has a square shape, thus:

$$H_{\text{sen}}(\xi, \eta) = \mathcal{F}\left(\frac{1}{c^2} \mathbb{1}_{|x|<c/2} * \mathbb{1}_{|y|<c/2}\right) e^{-\beta_1 c|\xi| - \beta_2 c|\eta|} = \text{sinc}(\xi c/2) \text{sinc}(\eta c/2) e^{-\beta_1 c|\xi| - \beta_2 c|\eta|} \quad (1.18)$$

where c represents the size of the sensor. The exponential term, takes into account the conductivity between one sensor and its neighbour. The second term in equation (1.17) takes into account the fact that the system might be moving during capture. If it moves a distance d at direction \mathbf{v} then:

$$H_{\text{mov}}(\mathbf{w}) = \mathcal{F}\left(\frac{1}{d} \mathbb{1}_{|\langle \mathbf{x}, \mathbf{v} \rangle| < d/2}\right) = \text{sinc}(\langle \mathbf{w}, \mathbf{v} \rangle d/2). \quad (1.19)$$

Finally, the optical system itself is not perfect and it also convolves the signal by a transfer function which might be very complicated. The following model is commonly assumed

$$H_{\text{opt}}(\xi, \eta) = e^{-\alpha c \sqrt{\xi^2 + \eta^2}} \quad (1.20)$$

for values of α of about 0.3 in our case, which means that in the useful region the image is low-pass filtered much more by the sensors themselves than by the optical system. Figure 1.6 shows some level lines of H for three instances of the previously described imaging system, where the (square) sensors of size c are assumed to be distributed at regular spacings of c on a square grid: (a) ccd1: an ideal case with a still camera ($d = 0$) where the optics and sensor isolation are assumed to be perfect ($\alpha = \beta_i = 0$); (b) ccd2: a less ideal case, with a still camera, but more realistic values for optics and conductivity ($\alpha \approx 0.30$ and $\beta_2 \approx 0.14$); (c) same situation as before but with the camera moving a distance $d = 2.35c$ along a direction \mathbf{v} which forms an angle of 11° with the horizontal.

Observe that all these systems produce highly aliased images. In figure 1.6(a) for instance, the transfer function only decayed from 1 to 0.6 at the border of the reciprocal cell. In order to reduce aliasing to a minimum we would ideally want the reciprocal cell to cover a larger region, e.g. up to the first zero-crossing of the sinc function. But that would require to double the sampling rate in both directions, which is not easy to

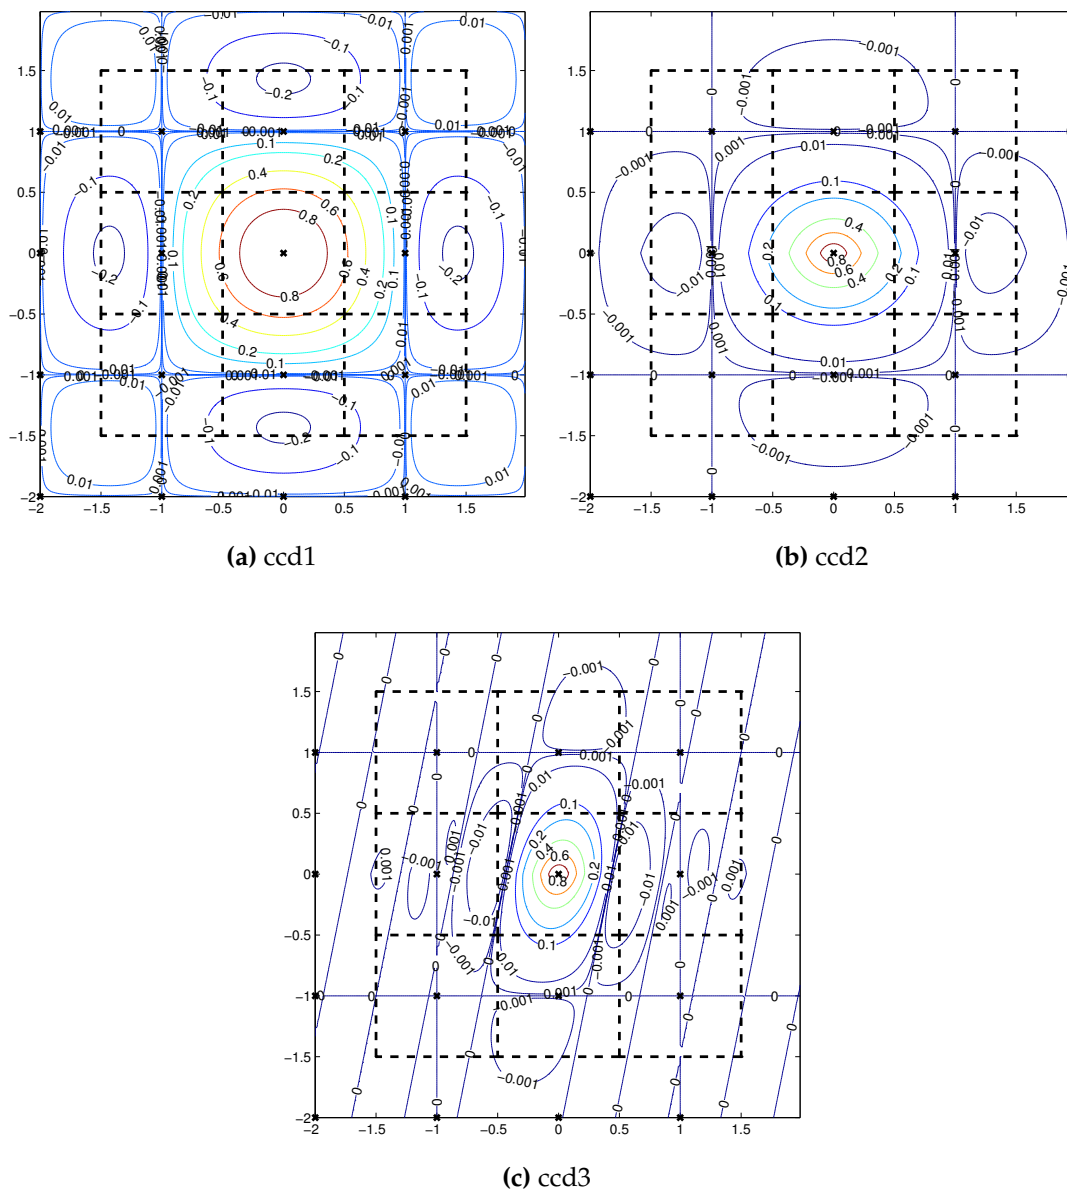


Figure 1.6: Modulation Transfer Functions and the corresponding Voronoi cells for three imaging systems composed of square CCD arrays. (a) Perfect optics and fixed camera; (b) more realistic optics; (c) the camera moves at uniform speed during capture.

Legend: In all subfigures the dashed lines represent the boundaries of the Voronoi reciprocal cell and its aliases. The X-shaped dots in the midpoint of each Voronoi cell represent the points in the dual (“aliasing”) grid Γ^* . The remaining curves are the level-lines of the MTF.

do without changing the size of the sensors, because this would require the sensors to overlap.⁵

However, several mechanisms have been proposed at the French Spatial Agency (CNES), which allow to increase the sampling rate without reducing the size of the sensors. Some of them (called hypermode and supermode) have been implemented on the SPOT5 earth observation satellite (launched in April this year) [116, 147]. Other novel mechanisms for improving the sampling rate involve both square and hexagonal grids. The corresponding transfer functions for all these systems are shown in figures 1.7, 1.8 and 1.9. More details about how such transfer functions are obtained from the real systems can be found in annex B, section 2.

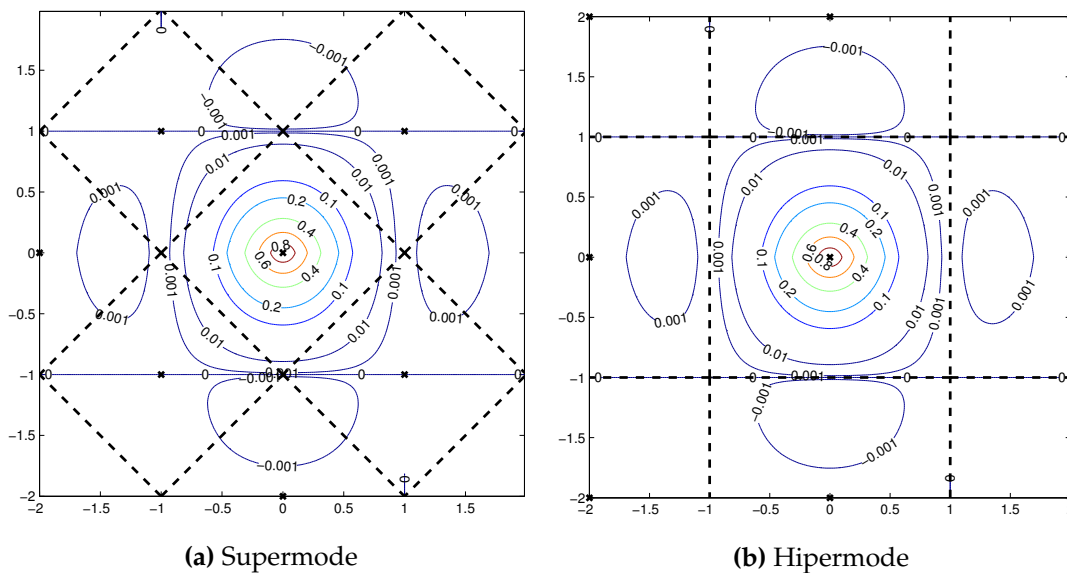


Figure 1.7: Modulation transfer function and the corresponding Voronoi cell for the Supermode sampling system used by the SPOT5 HRG earth observation satellite, and the Hipermode sampling system which has also been considered during its design. (See also the legend in figure 1.6.)

None of these systems exactly satisfies Shannon's compact support condition, and therefore they produce aliased images. But for some of them (see for instance figure 1.9(c)) aliasing effects may be far below the noise level and thus imperceptible, meaning that we unnecessarily increased the sampling rate. In order to find the right balance between increasing resolution and reducing aliasing artifacts on one hand, and

⁵ On the other hand reducing the size of the sensors wouldn't help, because this zero crossing would move to a higher frequency, and we would have the same aliasing problem, just at another scale.




Figure 1.8: Modulation transfer function and the corresponding Voronoi cell for some high-resolution square sampling systems. (See also the legend in figure 1.6.)

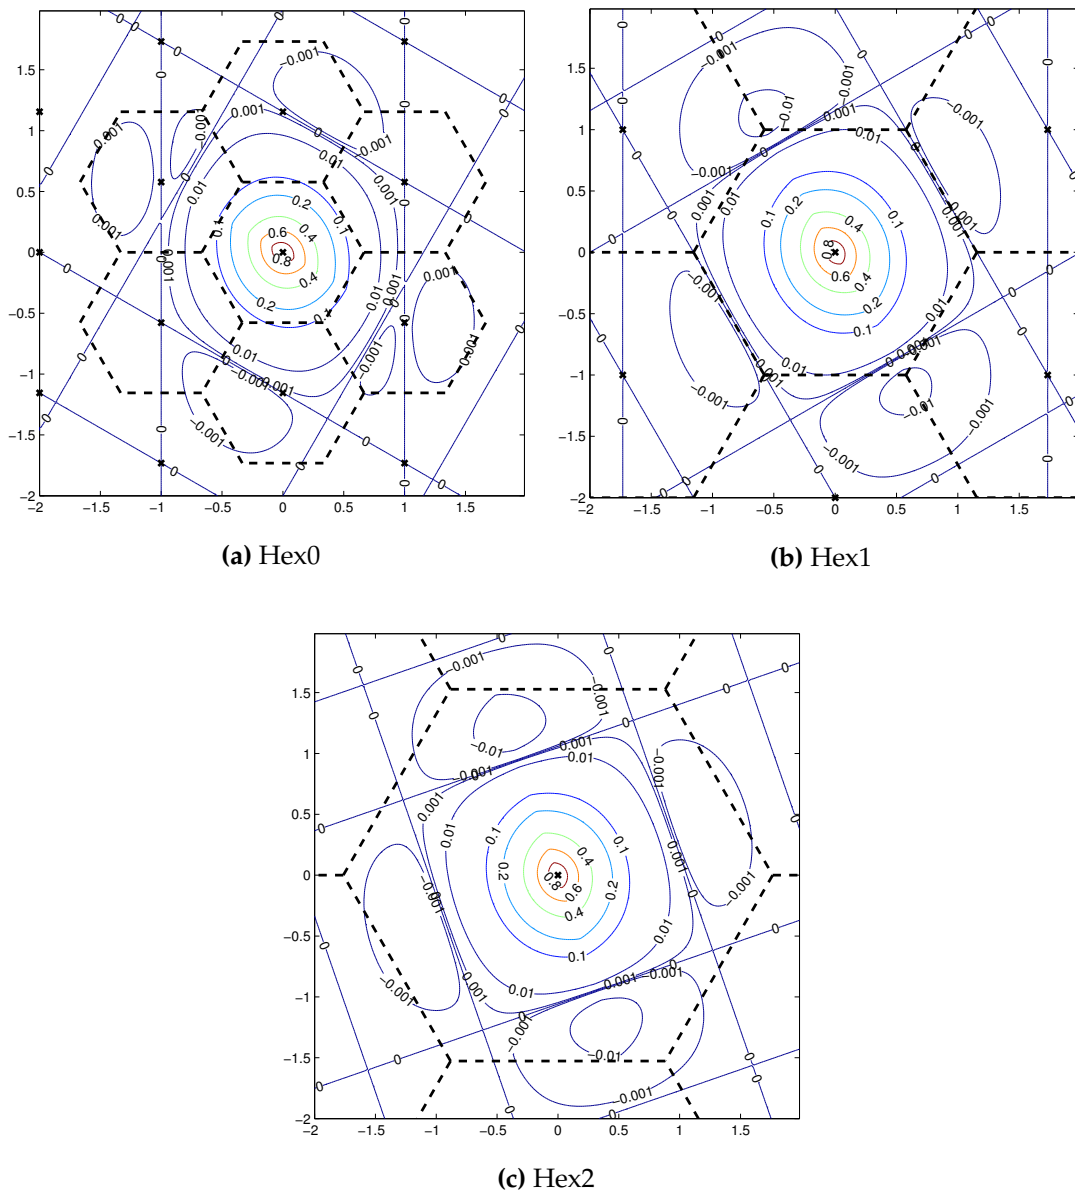


Figure 1.9: Modulation transfer function and the corresponding Voronoi cell for some high-resolution hexagonal sampling systems. (See also the legend in figure 1.6.)

on the other reducing the sampling rate to a minimum (which is dictated by bandwidth limitations), we need to extend the notions of critical sampling, sampling density and sampling efficiency introduced in section 2.4 to the case of infinite spectral support.

In this sense, it has been proposed [147] to consider the *essential spectral support* of the image to be sampled, namely:

$$\text{ess supp}(\hat{f}) := \left\{ \xi \in \mathbb{R}^2 : |\hat{f}| \geq \sigma|\Omega| \right\}. \quad (1.21)$$

i.e. the region over which the spectrum of the image is larger than the noise level (here we assumed a white iid Gaussian noise with variance σ^2 on an image domain Ω). Then we can apply the usual critical sampling and sampling efficiency definitions with the essential instead of the real spectral support. In [67] this criterion has been refined in order to take into account the signal to noise ratio at each location within the essential spectral support. The result is a resolution measure which is invariant with respect to any linear post-processing that can be done to reduce noise or aliasing. In the next section we further develop these ideas, and propose a resolution measure which will also allow us to choose the best reciprocal cell (section 5).

4 Resolution and information measures

4.1 Nominal Resolution

Usually resolution is defined in terms of dots-per-inch (dpi), lines-per-inch (lpi), etc. which measures the sampling density. This usual definition of resolution is what we shall call *spectral nominal resolution* $\sqrt{r_{\text{nom}}}$ (the reason for the word spectral will become clear in a minute), which we define precisely as the density of the sampling grid:

$$r_{\text{nom}} := \rho(\Gamma)$$

Note that r_{nom} is measured in terms of sampling points per unit area (*e.g.* dots per square inch). To obtain a dpi-style measure of resolution we simply use $\sqrt{r_{\text{nom}}}$. We also define a *spatial nominal resolution* which is proportional to the distance between neighboring samples, namely:

$$R_{\text{nom}} = r_{\text{nom}}^{-\frac{1}{2}}.$$

Usually we will measure these resolutions in terms of c where c^2 is the area covered by each of the sensors. Thus:

- a typical square CCD array (Sq0) has nominal resolution $r_{\text{nom}} = 1c^{-2}$, $R_{\text{nom}} = 1c$.

- a pair of CCD arrays in quincunx arrangement (SPOT5's supermode) has nominal resolution $r_{\text{nom}} = 2c^{-2}$, $R_{\text{nom}} = \frac{1}{\sqrt{2}}c$.
- a system of four CCD arrays (Hipermode) has nominal resolution $r_{\text{nom}} = 4c^{-2}$, $R_{\text{nom}} = \frac{1}{2}c$
- the hexagonal system Hex1 has $r_{\text{nom}} = 2\sqrt{3}c^{-2} \approx 3.46c^{-2}$, $R_{\text{nom}} = \frac{1}{\sqrt{2\sqrt{3}}}c$

Now consider a tile D^* of the dual grid Γ^* , for instance the Voronoi cell of Γ^* . Given the relationship between Γ and Γ^* we can show that the spectral nominal resolution

$$r_{\text{nom}} = \frac{1}{(2\pi)^2} \frac{1}{\rho(\Gamma^*)} = \frac{1}{(2\pi)^2} \int_{D^*} 1 \quad (1.22)$$

is in fact proportional to the area of the Voronoi cell of the dual grid Γ^* , whereas the spatial nominal resolution R_{nom}^2 is proportional to (the square root of) the area of the Voronoi cell of the sampling grid Γ (which explains the names we chose).

Considering noise and aliasing in the resolution measure. The spectrum of the sampled image \hat{g} may contain at different points very different amounts of noise and aliasing, which is not reflected by equation (1.22) which gives the same weight 1 to all Fourier coefficients of \hat{g} .

Assume an image acquisition system is defined by a triplet (Γ, H, n) . Then, according to equation (1.16) the Fourier transform of the acquired image is:

$$\begin{aligned} \hat{g} &= \mathcal{F}(\Delta_\Gamma) * (H \cdot \hat{f}) + \hat{n} \\ &= |D^*| \Delta_{\Gamma^*} * (H \cdot \hat{f}) + \hat{n} \\ &= |D^*| \sum_{\omega_0 \in \Gamma^*} (H \cdot \hat{f})(\cdot + \omega_0) + \hat{n} \end{aligned} \quad (1.23)$$

which we can rewrite in terms of

$$G = |D^*| H \cdot \hat{f} \quad (1.24)$$

$$G_{\text{alias}}(\xi) = \sum_{\gamma^* \in \Gamma^* \setminus \{0\}} G(\xi + \gamma^*) \quad (1.25)$$

as follows:

$$\hat{g} = G + G_{\text{alias}} + \hat{n}. \quad (1.26)$$

This last expression shows clearly the aliasing and noise errors present in the sampled image g . Thus it is not fair to consider all pixels equal. In the next section we shall refine our resolution measure by substituting the 1 in equation (1.22) by a weighting function $w(G(\xi), G_{\text{alias}}(\xi), \hat{n}(\xi))$, which is close to 1 in the absence of aliasing and noise, and close to 0 when either aliasing and noise dominates over the signal.

4.2 Effective resolution

Once such a weighting function w is defined we can define the effective resolution as:

$$r_{\text{eff}}(\Gamma, D^*, H, n, f) := \frac{1}{(2\pi)^2} \int_{D^*} w(G(\xi), G_{\text{alias}}(\xi), \hat{n}(\xi)) d\xi. \quad (1.27)$$

But this definition has two problems. First, it depends on the image f *before sampling* which is unknown. And secondly it will give a different value for each realization of the image f and the noise n . We would be tempted to solve both problems by defining

$$r_{\text{eff}}(\Gamma, D^*, H, \mathcal{N}, \mathcal{I}) := \mathbf{E}\{r_{\text{eff}}(\Gamma, D^*, H, n, f)\} \quad (1.28)$$

Where the expected value is taken with respect to the random variables $n \sim \mathcal{N}$ and $f \sim \mathcal{I}$. We shall come back later to the problems posed by the computation of such expected value. For the moment let us concentrate on the choice of a suitable weighting function w .

Conditions on the weighting function w . A first condition we want to impose is that w be independent of any invertible⁶ linear filtering k that we can perform on the digitized image g . The reason for this is that it is very common to apply “anti-aliasing” or “denoising” filters to an image after sampling, in order to reduce these artifacts (in section 6 we give an example of such manipulations). However, we want our measure to reflect the resolution of the “best image” that we can obtain by such manipulations from the acquisition system.

In order to express this condition in terms of w observe that $\mathcal{F}(k * g) = \hat{k} \cdot G + \hat{k} \cdot G_{\text{alias}} + \hat{k} \cdot \hat{n} = \hat{k} \cdot \hat{g}$. Hence we require that:

$$w(\hat{k} \cdot G, \hat{k} \cdot G_{\text{alias}}, \hat{k} \cdot \hat{n}) = w(G, G_{\text{alias}}, \hat{n}), \text{ for all } k \quad (1.29)$$

or put another way, the weighting function only depends on *relative* noise and aliasing.

$$w(G, G_{\text{alias}}, \hat{n}) = w\left(1, \frac{G_{\text{alias}}}{G}, \frac{\hat{n}}{G}\right) = W\left(\frac{G_{\text{alias}}}{G}, \frac{\hat{n}}{G}\right). \quad (W0)$$

Next, we don't care about the *phase* of relative noise or aliasing, we just consider its *magnitude*.

$$W(a, b) = W(|a|, |b|). \quad (W1)$$

⁶ Here we mean that k should be invertible in the sense that $\hat{k} \neq 0$ almost everywhere.

Finally, we should impose a few conditions, to ensure that W actually behaves like a weighting function, which penalizes high levels of relative noise and aliasing:

$$W(a, b) \in [0, 1] \quad (\text{W2})$$

$$W(0, 0) = 1 \quad (\text{W3})$$

$$W(+\infty, b) = W(a, +\infty) = 0 \quad (\text{W4})$$

$$W(\cdot, b) \text{ non-increasing for all } b, W(a, \cdot) \text{ non-increasing for all } a. \quad (\text{W5})$$

This still leaves a huge number of possibilities. For our experiments we considered the following two functions:

$$W(a, b) = \left(1 - \frac{a}{\theta_{\text{alias}}}\right)^+ \left(1 - \frac{b}{\theta_{\text{noise}}}\right)^+, \quad (1.30)$$

$$W(a, b) = \min\left\{\left(1 - \frac{a}{\theta_{\text{alias}}}\right)^+, \left(1 - \frac{b}{\theta_{\text{noise}}}\right)^+\right\}. \quad (1.31)$$

Observe that both functions satisfy a stronger condition than equation (W4). Namely, $W(a, b) = 0$ whenever $a \geq \theta_{\text{alias}}$ or $b \geq \theta_{\text{noise}}$, *i.e.* when either the relative aliasing or noise exceeds the specified thresholds. In principle we should set $\theta_{\text{alias}} \approx \theta_{\text{noise}} \approx 1$ meaning that whenever the aliasing or the noise are larger than the real signal, then we should ignore the corresponding Fourier coefficient because it contains almost no useful information. However, as we shall see in section 6, our experiments indicate that we obtain better results for $\theta_{\text{alias}} \approx 0.2$ and $\theta_{\text{noise}} \approx 5$, which may indicate that our perceptual system tolerates noise much better than aliasing artifacts.

Computing the expected value in equation (1.28). In order to compute the expected value in equation (1.28) we need suitable noise and image models. The noise model doesn't pose any serious problems, usually a white noise model $\mathcal{N} = N(0, \sigma)$ is accurate enough. But accurate statistical image models \mathcal{I} are unknown in general. In its absence we can compute the empirical expected value by taking the mean of equation (1.27) for a large number of images. But this is computationally very expensive, and its mathematical treatment is also much too complicated!!

A more sensible choice is to modify equation (1.28), by evaluating w not on a particular realization of the relative aliasing and noise, but on its expected value directly:

$$r_{\text{eff}}(\Gamma, D^*, H, \mathcal{N}, \mathcal{I}) := \frac{1}{(2\pi)^2} \int_{D^*} W \left(\sqrt{\frac{\mathbb{E}\{|G_{\text{alias}}(\xi)|^2\}}{\mathbb{E}\{|G(\xi)|^2\}}}, \sqrt{\frac{\mathbb{E}\{|\hat{n}(\xi)|^2\}}{\mathbb{E}\{|G(\xi)|^2\}}} \right) d\xi \quad (1.32)$$

Note that this expression is much simpler, since it doesn't take into account any possible dependency between $\hat{f}(\xi)$ and $\hat{f}(\xi')$ for $\xi \neq \xi'$, except for the term $\mathbb{E}\{|G_{\text{alias}}(\xi)|\}$

which considers the dependency between $\hat{f}(\xi)$ and its aliases $\hat{f}(\xi + \gamma^*)$ for all $\gamma^* \in \Gamma^*$. We will also simplify this term by substituting $\sqrt{\mathbb{E}\{|(H\hat{f})_{\text{alias}}(\xi)|^2\}}$ by $\sqrt{|HF|_{\text{alias}}^2(\xi)}$ where $F(\xi) = \sqrt{\mathbb{E}\{|\hat{f}(\xi)|^2\}}$. So, we define:

$$r_{\text{eff}}(\Gamma, D^*, H, N, F) := \frac{1}{(2\pi)^2} \int_{D^*} W \left(\sqrt{\frac{|HF|_{\text{alias}}^2(\xi)}{|HF|^2(\xi)}}, \sqrt{\frac{N^2(\xi)}{|HF|^2(\xi)}} \right) d\xi \quad (1.33)$$

which no longer depends on the image and noise models \mathcal{I}, \mathcal{N} , but just on (the square root of) the second order moments $F(\xi) = \sqrt{\mathbb{E}\{|\hat{f}(\xi)|^2\}}$ and $N(\xi) = \sqrt{\mathbb{E}\{|\hat{n}(\xi)|^2\}}$, for $n \sim \mathcal{N}, f \sim \mathcal{I}$. Now for a white noise model $\mathcal{N} = N(0, \sigma)$, we just have a constant $N(\xi) = \sigma$. (Note however that this study works equally well for any other zero-mean colored Gaussian noise model, which is determined by the function $N(\xi)$). Finally if the images are supposed to be in the BV space (bounded variation), a reasonable assumption is $F(\xi) \sim \frac{1}{|\xi|}$ except near the origin. In fact this assumption is quite accurate: by taking the empirical expected value $\mathbb{E}\{|\hat{f}(\xi)|^2\}$ for a large family of images, we obtain a function $F(\xi)$ whose shape is quite similar to that of $\frac{1}{|\xi|}$ (see figure 1.10).




Figure 1.10: Learning the decay rate of Fourier coefficients of natural images. After computing (the square root of) the mean of $|\hat{f}|^2$ for several natural images f , we obtain a roughly rotationally invariant function F , which was used as image model in our experiments. Here the mean was taken over each ring where $r \leq |\xi| < r + 1$. The result shows that the Fourier coefficients decay like $|\hat{f}(\xi)| \approx C|\xi|^{-p}$ with $1 \leq p \leq 2$, and typically $p \approx 1.6$.

We note that despite the many simplifications we did between equations (1.28) and (1.33), we find that both expressions yield experimentally comparable r_{eff} values, which indicates that these simplifications did not essentially affect the notion of resolution that we wanted to capture initially. We can conclude that the image model we used in equa-

tion (1.33), namely a coloured Gaussian noise, where $|\hat{f}(\xi)| \sim N(0, F(\xi)^2)$, is accurate enough for computing image resolution as defined in equation (1.28).

4.3 Information density.

The previous definition of effective resolution was motivated by perceptual arguments on the influence of noise and aliasing artifacts, that lead us to a certain choice of weighting function W .

On the other hand we can attempt to measure the amount of information of the original image $h * f$ that is preserved by the samples g . So in this section we will choose as weighting function the mutual information between the Fourier coefficient $\hat{g}(\xi)$ of the sampled image and the corresponding Fourier coefficient of the original image $G(\xi) = H(\xi)\hat{f}(\xi)$. This choice leads to a measure of effective information density

$$I_{\text{eff}} := \frac{1}{(2\pi)^2} \int_{D^*} I(\hat{g}(\xi), G(\xi)) d\xi \quad (1.34)$$

which measures the number of bits of information (per unit area) of the original image that can be obtained from the sampled image. The reciprocal cell will be chosen, as in the previous section in such a way that (Γ^*, D^*) is a packing.

Observe that this packing condition means that, if all $\hat{f}(\xi)$ are independent then all $G(\xi)$ and $\hat{g}(\xi')$ are independent except when $\xi = \xi'$. Hence the integral $\int_{D^*} I(\hat{g}(\xi), G(\xi)) d\xi$ equals the mutual information $I(\hat{g}, G|_{D^*})$ between the sampled image (in the Fourier domain) and the Fourier transform of the original image $G = H\hat{f}$ restricted to the reciprocal cell D^* .

We could still refine this model by computing $I(\hat{g}, G)$ without restriction to a reciprocal cell, and without assuming the independence between the Fourier coefficients of \hat{f} , but for the moment we are interested in a simpler model which still allows some qualitative comparisons with our results of the previous section on resolution measures.

In any case the mutual information $I(\hat{g}(\xi), G(\xi))$ inside the integral (and hence the information density measure I_{eff}) also depends on the statistical model we choose for the image. The following theorem gives a general lower bound, as well as the value of I_{eff} for two simple image models. In order to simplify the notation we express the

theorem in terms of the following dimensionless variables:

$$a^2(\xi) = \frac{\text{Var}\{G_{\text{alias}}(\xi)\}}{\text{Var}\{G(\xi)\}} \quad (\text{relative aliasing}) \quad (1.35)$$

$$b^2(\xi) = \frac{\text{Var}\{\hat{n}(\xi)\}}{\text{Var}\{G(\xi)\}} = \frac{|N|^2(\xi)}{|HF|^2(\xi)} \quad (\text{relative noise}) \quad (1.36)$$

$$k^2(\xi) = \frac{\text{Cov}\{G(\xi), G_{\text{alias}}(\xi)\}^2}{\text{Var}\{G(\xi)\}\text{Var}\{G_{\text{alias}}(\xi)\}} \quad (\text{relative covariance}) . \quad (1.37)$$

Observe that in the case where $\hat{f}(\xi)$ is independent of $\hat{f}(\xi')$ for all $\xi \neq \xi'$, we have a more convenient way of writing the relative aliasing, namely:

$$a^2(\xi) = \frac{|HF|_{\text{alias}}^2(\xi)}{|HF|^2(\xi)} \quad (1.38)$$

where $F(\xi) = \sqrt{\text{E}\{\hat{f}(\xi)\}}$ as in the previous section. Next we simply state the theorem. The details of the proof are given in appendix B

Theorem 2. *Given a sampled image g which relates to the original image $h * f$ through equations (1.16) and (1.26), the mutual information between $\hat{g}(\xi)$ and $G(\xi) = H(\xi)\hat{f}(\xi)$ is always lower bounded by:⁷*

$$I(\hat{g}(\xi), G(\xi)) \geq \frac{1}{2} \log_2 \left(\frac{2^{2\text{H}(G(\xi))} + 2^{2\text{H}(G_{\text{alias}}(\xi))} + 2^{2\text{H}(\hat{n}(\xi))}}{2^{2\text{H}(G_{\text{alias}}(\xi) + \hat{n}(\xi))}} \right). \quad (1.39)$$

This lower bound is achieved when $G(\xi)$, $G_{\text{alias}}(\xi)$ and $\hat{n}(\xi)$ are all independent zero-mean Gaussian variables. In that case the mutual information reduces to:

$$I(\hat{g}(\xi), G(\xi)) = \frac{1}{2} \log_2 \left(1 + \frac{1}{a^2 + b^2}(\xi) \right). \quad (1.40)$$

Finally, if we refine the previous model, by considering the dependence between $G(\xi)$ and $G_{\text{alias}}(\xi)$ through its covariance, the mutual information is:

$$I(\hat{g}(\xi), G(\xi)) = \frac{1}{2} \log_2 \left(1 + \frac{1 + a^2 k^2 + 2ak}{a^2(1 - k^2) + b^2}(\xi) \right). \quad (1.41)$$

Non-independent Gaussian Modeling. Observe that in the last equation of the theorem, we considered the dependence between $G(\xi)$ and $G_{\text{alias}}(\xi)$, since both terms depend on the Fourier coefficients $\hat{f}(\xi + \gamma^*)$ (for $\gamma^* \in \Gamma^*$) of the original image, which are not necessarily independent. On the other hand we still consider $\hat{n}(\xi)$ to be independent of the other terms since it is usually a white noise.

⁷ Note that here H denotes the entropy operator, whereas H denotes the modulation transfer function.

Effective and Nominal Information Density. Finally our information density measure becomes, in the non-independent Gaussian case:

$$I_{\text{eff}}(\Gamma, D^*, H, N, F, k) = \frac{1}{(2\pi)^2} \int_{D^*} \frac{1}{2} \log_2 \left(1 + \frac{1 + a^2 k^2 + 2ak}{a^2(1 - k^2) + b^2}(\xi) \right) d\xi \quad (1.42)$$

which can be written as

$$I_{\text{eff}} = \frac{1}{(2\pi)^2} \int_{D^*} W_k(a(\xi), b(\xi)) d\xi \quad (1.43)$$

with $W_k(a, b) = \frac{1}{2} \log \left(1 + \frac{1 + a^2 k^2 + 2ak}{a^2(1 - k^2) + b^2} \right)$. Observe that if we consider aliasing to be independent from the original image ($k = 0$), then $W_0(a, b) = \frac{1}{2} \log \left(1 + \frac{1}{a^2 + b^2} \right)$. We observe that this function W_0 satisfies all requirements (W0-W5) except for (W2) and (W3). Instead $W_0(a, b) \in [0, +\infty]$ with $W_0(0, 0) = +\infty$. But this is normal since we are not measuring resolution in terms of sampling per unit area, we are measuring information in terms of bits per unit area.

For $k \neq 0$, W_k still satisfies (W0,W1,W5) and it partially satisfies (W4). In fact $W_k(a, +\infty) = 0$ but $W_k(+\infty, b)$ may be positive. This happens for instance when $\text{Var}\{G(\xi)\} \ll \text{Var}\{G_{\text{alias}}(\xi)\}$. In this case we have infinite relative aliasing, but the sampled coefficient (essentially $G_{\text{alias}}(\xi)$) still gives some information about $G(\xi)$ through the covariance between $G(\xi)$ and $G_{\text{alias}}(\xi)$.

An ideal system with no aliasing ($H = \mathbb{1}_{D^*}$) and no noise $N(\xi) = 0$ would yield infinite information density, but as a reference to compare with our system, we can consider the “nominal information density” to be the one that is obtained by a semi-ideal system with no aliasing ($H = \mathbb{1}_{D^*}$) and a minimal amount of noise $N(\xi) = \varepsilon$.

$$I_{\text{nom}}(\Gamma, D^*, \varepsilon, F) := \frac{1}{(2\pi)^2} \int_{D^*} \frac{1}{2} \log_2 \left(1 + \frac{|F(\xi)|^2}{\varepsilon^2} \right) d\xi \quad (1.44)$$

Aliasing vs. Noise in Information and Resolution Measures. The previous definition of effective information density shows that for a given level of relative distortion d , if this distortion comes from (image-independent) noise, then information decreases to a larger extent than if it comes from (image-dependent) aliasing. In fact, in the first case, $a = 0, b = d$ and $W(a, b) = \frac{1}{2} \log_2 \left(1 + \frac{1}{d^2} \right)$, whereas in the second case $a = d, b = 0$ and $W(a, b) = \frac{1}{2} \log_2 \left(1 + \frac{1 + k^2 d^2 + kd}{1 - k^2} \right)$, which is larger.

With respect to the resolution measures introduced in the previous subsections, we observed the contrary: for a given level of relative distortion, if it comes from noise, resolution decreases less than if it comes from aliasing (cf. the discussion on the choice of parameters $\theta_{\text{noise}} \gg \theta_{\text{alias}}$).

This comes from the fact that image-independent distortions such as white noise are more easily ignored by our perceptual system than image-dependent distortions such as aliasing. From an information-theoretical point of view, however, the latter gives us more information than the former for reconstructing the original image, provided we know the statistics of the dependence between the distortion (aliasing) and the original image.

So, if we know how to correct aliasing artifacts from knowledge of image statistics, then it makes sense to introduce this statistical knowledge in our information measure. Otherwise, it is better to consider it as noise and eliminate it.

5 Optimal Reciprocal Cell

In the previous section, all measures of image resolution or information that we defined, depend upon the shape of the reciprocal cell D^* . In this section we address the problem of choosing such a reciprocal cell. However, once we defined our measures of image resolution and information, we just define the optimal reciprocal cells as those that maximize either resolution or information. This presents, however some technical difficulties, so we shall first introduce a general technique to construct tilings, and then we shall apply it to our information and resolution measures.

We start with the definition of the canonical reciprocal cell associated to a resolution measure u and an aliasing grid Γ^* .

Definition 5 (Canonical Tiling). *Given Γ^* a regular grid in \mathbb{R}^d , and a function $u : \mathbb{R}^d \rightarrow \mathbb{R}$, the set*

$$R_{\max} = \{\xi \in \mathbb{R}^d : u(\xi) > u(\xi + \gamma) \text{ for all } \gamma \in \Gamma^*, \gamma \neq 0\} \quad (1.45)$$

is called canonical tiling associated to the function and the grid.

Since we are selecting among all possible aliases of a frequency ξ , the one which gives the highest resolution $u(\xi + \gamma)$, it is quite clear that if R_{\max} is a tiling, then it maximizes $\int_R u(x)dx$ with respect to all packings R .

It is also clear that R_{\max} is a packing because we only select one point within the set $\{\xi + \gamma\}_{\gamma \in \Gamma^*}$.

Showing that the “canonical tiling” is actually a tiling and the only one which maximizes $\int_R u(x)dx$ is a little bit more technical, and requires some further properties that we state in the following definition.

Definition 6. Given a regular grid $\Gamma^* \subset \mathbb{R}^d$, a function $u : \mathbb{R}^d \rightarrow \mathbb{R}$ is called an admissible alias selector if it satisfies the following conditions

(A1) The associated canonical tiling R_{\max} is bounded.

(A2) For almost every $\xi \in \mathbb{R}^d$, there is a single $\gamma_0 \in \Gamma^*$ such that $u(\xi + \gamma_0) = \sup_{\gamma \in \Gamma^*} u(\xi + \gamma)$.

The purpose of condition (A1) is to provide reciprocal cells $\overline{R_{\max}}$ with compact support which allow us to apply Shannon's sampling theorem. The purpose of condition (A2) is to avoid some situations that might lead R_{\max} to be a sub-tiling. In the sequel we will use the following (equivalent but more convenient) formulation for condition (A2):

(A2') The associated anomaly set A_{u, Γ^*} defined below has zero measure.

$$A_{u, \Gamma^*} = \{ \xi \in \mathbb{R}^d : \text{card}(G_{\xi, u, \Gamma^*}) \neq 1 \}$$

$$\text{where } G_{\xi, u, \Gamma^*} = \left\{ \gamma_0 \in \Gamma^* : u(\xi + \gamma_0) = \sup_{\gamma \in \Gamma^*} u(\xi + \gamma) \right\}.$$

The examples below illustrate how we might fail to meet this condition.

Example. If u is Γ -periodic then R_{\max} is clearly empty. But an admissible alias selector cannot be Γ -periodic, since periodicity would imply that $G_{\xi, u, \Gamma^*} = \xi + \Gamma^*$ for any ξ . Thus the anomaly set would be all of \mathbb{R}^d .

Example. If $|u(x)|$ monotonically increases to u_{\max} when $|x| \rightarrow \infty$ without ever attaining this value, then R_{\max} is also empty. But an admissible alias selector will never show such a behaviour, since it would imply that $G_{\xi, u, \Gamma^*} = \emptyset$ for any ξ . Thus the anomaly set would be all of \mathbb{R}^d also in this case but for a different reason.

Appendix C gives some simpler sufficient conditions for a function u to be an admissible alias selector. However, the previous examples cover almost all anomalous cases.

Proposition 1. Let Γ^* be a regular grid in \mathbb{R}^d . For any function $u : \mathbb{R}^d \rightarrow \mathbb{R}$, the associated canonical tiling R_{\max} is such that (Γ^*, R_{\max}) is a packing of \mathbb{R}^d .

If in addition u is an admissible alias selector then (Γ^*, R_{\max}) is also a covering and hence a tiling.

Finally, such a tiling is the unique⁸ solution to the following optimization problem:

$$R_{\max} = \arg \sup_{\{R : (\Gamma^*, R) \text{ is a packing}\}} \int_R v(u(x)) dx \quad (1.46)$$

⁸ up to a set of zero measure

for any non-negative, strictly increasing function $v : [u_{\min}, \infty) \rightarrow [0, \infty)$ where $u_{\min} = \inf_x u(x)$.

Observe that if we do not require uniqueness of R in the sense of maximization of $\int_R v(u(x))dx$ we can avoid the technical hypothesis that u is an admissible alias selector by modifying the definition of R_{\max} as follows:

$$R_{\max} = \{ \xi \in \mathbb{R}^d : \text{for all } \gamma \in \Gamma^*, \gamma \neq 0 \implies u(\xi) > u(\xi + \gamma) \text{ or} \\ u(\xi) = u(\xi + \gamma) \text{ and } \xi \prec (\xi + \gamma) \}$$

where \prec is any arbitrary order in \mathbb{R}^d (e.g. compare first the magnitude, then the angle to the first axis, etc.). With this definition proposition 1 still holds (except for the uniqueness of R_{\max}), and we do not need the hypothesis that u be an admissible alias selector. Then instead of uniqueness, we resolved the ambiguity in the maximization by introducing an arbitrary ordering of the frequency space. Here we judged uniqueness as an important property, and it is easy to show that the functions we considered are admissible alias selectors. For this reason we keep the original definition of R_{\max} .

Proof. To prove that (Γ^*, R_{\max}) is a packing, consider $x \in R_{\max}$. Then for all $\gamma \in \Gamma^* \setminus \{0\}$ we have $u(x) > u(x + \gamma)$. This means that $(x + \gamma) \notin R_{\max}$, so $x \notin (R_{\max} + \gamma)$, and we conclude that (Γ^*, R_{\max}) is a packing.

To prove that (Γ^*, R_{\max}) is a covering consider any $x \in \mathbb{R}^d \setminus A_{u, \Gamma^*}$. Because $x \notin A_{u, \Gamma^*}$, we have that G_{x, u, Γ^*} consists of a single element γ_x and $(x + \gamma_x) \in R_{\max}$. Since A_{u, Γ^*} has zero measure this means that for almost every point $x \in \mathbb{R}^d$, there is $\gamma_x \in \Gamma^*$ such that $x \in (R_{\max} + \gamma)$, so (Γ^*, R_{\max}) is a covering.

To show the last part of the proposition, observe first that for all $x \in R_{\max}$, $v(u(x)) > 0$. In fact, since v is increasing, $v(u(x)) = 0$ would mean $u(x) = u_{\min}$, and consequently $u(x) \leq u(y)$ for all y , and in particular for $y = x + \gamma$, with $\gamma \in \Gamma^*$. Hence $v(u(x)) = 0$ implies $x \notin R_{\max}$, which proves that for $x \in R_{\max}$, $v(u(x)) > 0$.

Next consider any packing (Γ^*, R) , and any point $x \in R \setminus R_{\max}$. Since (Γ^*, R_{\max}) is a tiling there is some $\gamma_x \in \Gamma^*$ such that $(x + \gamma_x) \in R_{\max}$, except if $x \in N_1$ where N_1 has zero measure. Note that $\gamma_x \neq 0$ (otherwise $x \in R_{\max}$ would contradict our initial assumption) and the packing condition of R implies that $(x + \gamma_x) \notin R$ except for $x \in N_1 \cup N_2$, where N_2 has zero measure. If we call τ the mapping $x \mapsto (x + \gamma_x)$ then τ maps $(R \setminus R_{\max}) \setminus (N_1 \cup N_2)$ (which we shall abbreviate as R_1) into a subset of $(R_{\max} \setminus R)$. Now from the definition of R_{\max} we deduce that $u(\tau(x)) > u(x)$, and since v

is increasing $v(u(\tau(x))) > v(u(x))$. Therefore

$$\int_R v(u(x))dx - \int_{R_{\max}} v(u(x))dx = \int_{R \setminus R_{\max}} v(u(x))dx - \int_{R_{\max} \setminus R} v(u(x))dx \quad (1.47)$$

$$= \int_{(R \setminus R_{\max}) \setminus (N_1 \cup N_2)} v(u(x))dx - \int_{R_{\max} \setminus R} v(u(x))dx \quad (1.48)$$

$$\leq \int_{\tau((R \setminus R_{\max}) \setminus (N_1 \cup N_2))} v(u(y))dy - \int_{R_{\max} \setminus R} v(u(x))dx \quad (1.49)$$

$$\leq 0. \quad (1.50)$$

Since in equation (1.49) the integration domain $\tau(R_1) \subseteq R_{\max}$, and we have $v(u(y)) > 0$, so the first inequality is tight iff $\tau(R_1)$ has zero measure. From the construction of τ and the fact that N_1 and N_2 have zero measure, we conclude that $R \setminus R_{\max}$ has zero measure. Then the second inequality is tight only if $((R_{\max} \setminus R) \setminus \tau(R_1))$ has zero measure. Using the previous condition this means that $(R_{\max} \setminus R)$ has zero measure.

So R is optimal iff it coincides with R_{\max} at almost every point. \square

The easiest way to define a tiling of the spectral domain, which produces little aliasing, is the following. Consider $u(\xi) = \text{Var}\{G(\xi)\}$ and define

$$D_{\text{tiling}}^* = \{\xi : \text{Var}\{G(\xi)\} > \text{Var}\{G(\xi + \gamma^*)\} \text{ for all } \gamma^* \in \Gamma^*, \gamma^* \neq 0\}. \quad (1.51)$$

According to the previous proposition then, $(\Gamma^*, D_{\text{tiling}}^*)$ is a tiling of the spectral domain. In addition, by its construction, it is a tiling that concentrates as much of the energy of the original image G as possible. In fact it maximizes $\int_{D^*} \text{Var}\{G(\xi)\}d\xi$ among all packings (Γ^*, D^*) . This tiling will be useful in the sequel, when we consider the resolution-optimizing reciprocal cell.

Similarly, we can apply the proposition to $u(\xi) = \frac{1+a^2k^2+ak}{a^2(1-k^2)+b^2}(\xi)$ and $v(u) = \log_2(1+u)$ to define the information maximizing reciprocal cell as:

$$D_{\text{I-opt}}^* := \left\{ \xi : \frac{1+a^2k^2+ak}{a^2(1-k^2)+b^2}(\xi) > \frac{1+a^2k^2+ak}{a^2(1-k^2)+b^2}(\xi + \gamma^*) \text{ for all } \gamma^* \in \Gamma^*, \gamma^* \neq 0 \right\}. \quad (1.52)$$

This reciprocal cell is in fact, according to proposition 1, a tiling that maximizes the effective information density (for a given sampling grid, image and noise model), among all packings of the spectral domain:

$$D_{\text{I-opt}}^* = \arg \sup_{\{D^*: (\Gamma^*, D^*) \text{ is a packing}\}} I_{\text{eff}}(\Gamma, D^*, H, N, F, k). \quad (1.53)$$

Its definition differs from D_{tiling}^* , even in the Gaussian independent case where $k = 0$. Nevertheless, our experiments show that in practice both criteria produce almost identical reciprocal cells (see figures 1.11 to 1.14).

We could define the resolution-optimal reciprocal cell in a similar way, but then we would not obtain a unique solution. This is due to the fact that in our effective resolution measure we included a thresholding operation that may prevent the support of $v \circ u$ to form a covering with Γ . Hence we can no longer apply proposition 1 to $u(\xi) = W(a(\xi), b(\xi))$, $v = Id$.

However we can consider a limiting case:

$$D_{\text{R-opt}}^* := \{\xi : a(\xi) < \theta_{\text{alias}} \text{ and } b(\xi) < \theta_{\text{noise}}\} \quad (1.54)$$

and characterize all resolution-optimizing reciprocal cells in terms of $D_{\text{R-opt}}^*$ and D_{tiling}^* , as follows:

Proposition 2. Consider W in equations (1.30-1.31) and assume that $\theta_{\text{alias}} \leq 1$. Then the maximum of

$$\int_{D^*} W(a(\xi), b(\xi)) d\xi$$

among all packings (Γ^*, D^*) , is attained whenever $D_{\text{R-opt}}^* \subseteq D^* \subseteq D_{\text{tiling}}^*$. (Inclusions are in the weak, “almost everywhere” sense, i.e. up to a set of zero measure).

Proof. First we show that $D_{\text{R-opt}}^* \subseteq D_{\text{tiling}}^*$. In fact $\xi \in D_{\text{R-opt}}^*$ iff:

$$\begin{aligned} & \theta_{\text{alias}} > a(\xi) \text{ and } \theta_{\text{noise}} > b(\xi) \\ \Rightarrow & \text{Var}\{G(\xi)\} > \theta_{\text{alias}}^{-2} \text{Var}\{G_{\text{alias}}(\xi)\} \geq \text{Var}\{G_{\text{alias}}(\xi)\} \\ & \geq \text{Var}\{G(\xi + \gamma^*)\} \text{ for all } \gamma^* \in \Gamma^*, \gamma^* \neq 0 \\ \Rightarrow & \xi \in D_{\text{tiling}}^*. \end{aligned} \quad (1.55)$$

Now observe that $D_{\text{R-opt}}^* = \text{supp}(W(a(\xi), b(\xi)))$. This means that the (non-constrained) maximum of $\int_{D^*} W(a(\xi), b(\xi)) d\xi$ is attained for any $D^* \supseteq D_{\text{R-opt}}^*$. If in addition $D^* \subseteq D_{\text{tiling}}^*$ then (Γ^*, D^*) is a packing and D^* solves the (constrained) maximization problem. \square

From the proof it is clear that $(\Gamma^*, D_{\text{R-opt}}^*)$ is a packing, and it becomes a tiling only if $\theta_{\text{alias}} = 1$, θ_{noise} is sufficiently large and the maximal alias $\text{Var}\{G(\xi + \gamma_{\text{max}}^*)\}$ is larger than the squared sum of the remaining aliases $\sum_{\gamma^* \in \Gamma^* \setminus \{0, \gamma_{\text{max}}^*\}} \text{Var}\{G(\xi + \gamma^*)\}$. Thus in general it will not be a tiling and all reciprocal cells D^* such that $D_{\text{R-opt}}^* \subseteq D^* \subseteq D_{\text{tiling}}^*$ will

maximize r_{eff} among all possible packings. Thus we shall keep the minimal resolution-optimizing reciprocal cell as a reference.

Finally let's observe that if both the image model F and the transfer function H are isotropic, then G is also isotropic and it is easy to verify that in this case both the information-maximizing reciprocal cell $D_{\text{I-opt}}^*$ and D_{tiling}^* coincide with the Voronoi reciprocal cell D_{Vor}^* .

6 Examples and Applications

6.1 Optimal Reciprocal Cells

Figures 1.11, 1.12, 1.13 and 1.14 show the optimal reciprocal cells that were obtained for the different sampling systems that were presented in section 3. In all simulations we assumed a signal to noise ratio of 48 dB⁹ Both the first and second column show the boundary of D_{tiling}^* (in black) and the boundary of $D_{\text{R-opt}}^*$ (in white) for $\theta_{\text{alias}} = \theta_{\text{noise}} = 0$. The difference between these two columns is the background information that is showed: In the first column the gray-level represents the relative aliasing in log-scale $\log_{10} a(\xi)$, whereas the second column shows the relative noise in log scale $\log_{10} b(\xi)$.

Finally the third column shows the information-optimal reciprocal cell $D_{\text{I-opt}}^*$, in the case where the Fourier coefficients of the image $\hat{f}(\xi)$ are modeled as independent, zero-mean Gaussian variables. Even though the definition of $D_{\text{I-opt}}^*$ is not exactly the same as that of $D_{\text{R-opt}}^*$, the resulting regions are hardly distinguishable in all cases.

Observe how, even in a simple situation like ccd3 that may arise even with a digital camera, the optimal reciprocal cell might consist of more than one connected component.

6.2 Resolution and Information measures

Once we obtained the optimal reciprocal cells we can compute the corresponding effective resolution and information density measures. These are displayed in table 1.1, along with some further information which gives some indication of sampling effi-

⁹ To compute the SNR in dB we apply the usual formula $dB = 10 \log_{10}(\frac{\|f\|^2}{\|n\|^2})$. Then a SNR of 48 dB (more precisely 48.1648... dB) is obtained when the standard deviation of the noise is $\sigma = \frac{\|f\|}{256S}$, where S is the area of the image domain and $\frac{\|f\|}{S}$ is the normalized L^2 norm of the image, or its empirical standard deviation if we assume it with zero mean.





Figure 1.11: Optimal reciprocal cells for the three CCD sampling systems in figure 1.6.


Legend: In all figures above, the dashed lines represent the boundaries of the Voronoi reciprocal cell and its aliases. The crosses represent the points in the dual (“aliasing”) grid Γ^* . The remaining curves are the level-lines of the MTF.

In the first and second columns the thick black curve represents the boundary of the optimal tiling D_{tiling}^* , and the thick white line represents the boundary of the resolution-optimal reciprocal cell $D_{\text{R-opt}}^*$. In the last column the thick white line represents the boundary of the information-optimal reciprocal cell $D_{\text{I-opt}}^*$.

The gray-scale represents: in the first column the value of the relative aliasing $\log_{10}a(\xi)$ in logarithmic scale; in the second column the value of the relative noise $\log_{10}b(\xi)$ also in logarithmic scale; in the third column the value of the mutual information density $\log_2\left(1 + \frac{1}{a^2 + b(2)}(\xi)\right)$ measured in bits/ c^2 .



(a) Supermode



(b) Hipermode

Figure 1.12: Optimal reciprocal cells for the sampling systems in figure 1.7. (See also the legend in figure 1.11.)




Figure 1.13: Optimal reciprocal cells for the three Square sampling systems in figure 1.8. (See also the legend in figure 1.11.)




Figure 1.14: *Optimal reciprocal cells for the three hexagonal sampling systems in figure 1.9. (See also the legend in figure 1.11.)*

System	R_{nom} (c)	R_{eff} (c)	I_{nom} (bits/c ²)	I_{eff} (bits/c ²)	A (%)	B (%)	$\frac{R_{\text{nom}}}{R_{\text{eff}}}$ (%)	$I_{\text{eff}} R_{\text{nom}}^2$ (bits/sample)
ccd1	1.0000	1.2601	8.18	2.45	32.98	0.38	79.36	2.45
ccd2	1.0000	1.1703	6.86	3.48	23.58	1.23	85.44	3.48
ccd3	1.0000	1.1525	5.63	3.95	23.14	5.72	86.77	3.95
Sq0	1.0000	1.1463	6.66	3.91	20.81	1.57	87.24	3.91
Hex0	0.9306	1.0380	7.44	5.14	18.84	1.97	89.65	4.45
Supermode	0.7071	0.7843	9.51	8.76	11.91	9.78	90.16	4.38
Sq1	0.7071	0.7761	9.99	8.91	13.68	8.68	91.11	4.46
Hex1	0.5373	0.6716	11.40	11.35	3.81	18.37	80.00	3.28
Hipermode	0.5000	0.7093	10.30	10.27	2.42	18.31	70.50	2.57
Sq2	0.4472	0.6653	11.68	11.67	1.67	18.37	67.22	2.33
Hex2	0.3517	0.6626	11.82	11.82	0.00	18.24	53.09	1.46

Table 1.1: Quantitative evaluation of the sampling system $(H, n, \Gamma, D_{\text{R-opt}}^*)$ with $\theta_{\text{alias}} = \theta_{\text{noise}} = 1$, and for $(H, n, \Gamma, D_{\text{I-opt}}^*)$ with independent Gaussian model for the image Fourier coefficients, and signal to noise ratio of 48 dB.

System	r_{nom} (samples/c ²)	r_{eff} (samples/c ²)	I_{nom} (bits/c ²)	I_{eff} (bits/c ²)	A (%)	B (%)	$r_{\text{eff}}/r_{\text{nom}}$ (%)	$I_{\text{eff}}/r_{\text{nom}}$ (bits/sample)
ccd1	1.00	0.63	8.18	2.45	32.98	0.38	63	2.45
ccd2	1.00	0.73	6.86	3.48	23.58	1.23	73	3.48
ccd3	1.00	0.75	5.63	3.95	23.14	5.72	75	3.95
Sq0	1.00	0.76	6.66	3.91	20.81	1.57	76	3.91
Hex0	1.15	0.93	7.44	5.14	18.84	1.97	80	4.45
Supermode	2.00	1.63	9.51	8.76	11.91	9.78	81	4.38
Sq1	2.00	1.66	9.99	8.91	13.68	8.68	83	4.46
Hex1	3.46	2.22	11.40	11.35	3.81	18.37	64	3.28
Hipermode	4.00	1.99	10.30	10.27	2.42	18.31	49	2.57
Sq2	5.00	2.26	11.68	11.67	1.67	18.37	45	2.33
Hex2	8.08	2.28	11.82	11.82	0.00	18.24	28	1.46

Table 1.2: Same as table 1.1 with spectral resolution measures expressed in (samples per unit area), instead of spatial resolution (expressed in mean distance between samples).

ciency and quality of the different systems.

The first two columns show the spatial nominal and effective resolutions R_{nom} and R_{eff} , both expressed as multiples of the size c of the sensors. Observe that at equal sampling rate (equal R_{nom}), the first four systems have quite different effective resolution. In particular, the case of the CCD array which is blurred due to a moving camera (ccd3) performs much better than the ideal CCD array (ccd1) where the only low-pass filter is performed by the sensors themselves. This is due to the high amount of aliasing present in ccd1, which is not that much the case in ccd3. This can be observed in columns 4 and 5 which show the average amount of relative aliasing and noise: $A = \int_{D^*} a(\xi) d\xi / |D^*|$ and $B = \int_{D^*} b(\xi) d\xi / |D^*|$ respectively. The sampling efficiency in terms of resolution is measured by the ratio $R_{\text{nom}}/R_{\text{eff}}$ which expresses how much of the nominal resolution (sampling rate) is effectively attained by this system.

Column four shows the information density measure I_{eff} , expressed in terms of bits of mutual information per unit area, where the unit area c^2 is the surface of a single sensor. This is to be compared with the nominal information density I_{nom} as defined in equation (1.44). Note that unlike nominal resolution I_{nom} , I_{eff} takes into account the amount of noise and the decay rate of Fourier coefficients in the image. We can also compare I_{eff} to the information density one would obtain in an even more ideal situation where we do not only have an ideal transfer function $H = \mathbb{1}_{D_{\text{Vor}}^*}$, but also an ideal decay rate of Fourier coefficients, *i.e.* $F(\xi) = \text{constant}$. In such an ideal case the information density would be $I_{\text{ideal}} = \frac{|RC|}{(2\pi)^2} \frac{1}{2} \log_2(1 + SNR)$. Since in our case $SNR = 48$ dB, we would have $I_{\text{ideal}} = \frac{|RC|}{(2\pi)^2} 8$ bits = $R_{\text{nom}}^2 8$ bits, and we would at most have 8 bits of mutual information per sample. Hence another way of representing the efficiency of sampling is the product $I_{\text{eff}} R_{\text{nom}}^2$ (last column) which represents the number of bits of mutual information per sample and is bounded by 8.

Observe that both the effective resolution and information measures yield Supermode and Sq1 as the most effective sampling systems. Resolution may be further increased, but the cost in terms of sampling rate per increase in resolution will start growing faster beyond this point.

6.3 Application to Restoration

Finally we show an unexpected application of these results. We started with the definition of a resolution measure. To do so, we realized that we needed to adapt the reciprocal cell. Otherwise, our resolution measure would be under-estimated. Now, once we found the optimal reciprocal cell, we should use it for restoration as well. The

procedure is as follows:

From the discrete image g obtained by equation (1.16) we compute the discrete Fourier transform \hat{g} . In order to remove aliasing from this image without losing any good information, we need to oversample. We do so either by adding zeros between the samples g before the Fourier transform, or equivalently by periodically extending \hat{g} . Then we multiply this Γ^* -periodic \hat{g} by $\mathbb{1}_{D_{R\text{-opt}}^*}$ and take the inverse Fourier transform to obtain a zoomed version of g with some aliased or noisy coefficients removed and some aliased coefficients re-placed at their correct locations. The sampled image g and the corrected $\overline{\mathcal{F}}(\hat{g}\mathbb{1}_{D_{R\text{-opt}}^*})$ can be observed in figures 1.15(b) and 1.16(b) respectively.

This actually improves the image but in most cases the difference is hardly perceptible. In order to observe the full advantage of this procedure we need to deconvolve the transfer function. We shall do so by two novel procedures.

Wiener Filter for sampled images. The classical way for restoring images by means of linear filters is the Wiener Filter which is usually derived in a continuous setting, which does not take sampling into account:

$$g = h * f + n. \quad (1.56)$$

With this assumption the Wiener Filter restoration is defined as

$$u = k * g \quad \text{where} \quad k = \arg \min_k \mathbb{E}\{\|k * g - f\|^2\} \quad (1.57)$$


where the statistical models for f and n are usually colored Gaussian with zero mean and known variance as in sections 4.2 and 4.3. The solution is found to be

$$\hat{k}(\xi) = \frac{\overline{H}(\xi)}{|H(\xi)|^2(1 + b^2(\xi))} \quad (1.58)$$


where $b^2(\xi)$ is the relative noise as defined in equation (1.36). Note that to derive this result a crucial assumption is that $\hat{f}(\xi)$, $\hat{n}(\xi)$ and $\hat{f}(\xi')$ are all independent whenever $\xi \neq \xi'$, which is as discussed before not a very accurate image model. Note also that the denominator has to be understood as $|H(\xi)|^2 + \frac{\sigma^2}{|F(\xi)|^2}$ when $H(\xi)$ tends to zero, we only used the notation above for convenience).

Now this formulation doesn't take sampling into account and this is why aliasing doesn't come up in the final result. If we start from the complete image model with sampling:


$$g = \Delta_\Gamma \cdot (h * f) + n \quad (1.59)$$



(a) Original image and its Fourier spectrum



(b) Sampled image zoomed by zero-padding



(c) Deconvolved by Total Variation

Figure 1.15: Sampling with the Sq2 system and deconvolution by Total Variation minimization. Observe how the aliased second harmonic of the roof's texture explodes when deconvolved by a wrong Fourier coefficient of the transfer function.







Figure 1.16: Deconvolution of the image in figure 1.15 by Total Variation minimization with the data-fitting term integrated over the optimal reciprocal cell D_{R-opt}^* . Observe how the errors in the previous figure are reduced by placing aliasing-dominated regions at their correct locations in the Fourier spectrum.

(a) Optimal reciprocal cell D_{R-opt}^* (shown in transparent dark gray), optimal tiling D_{tiling}^* (shown in transparent medium gray), and Voronoi cell D_{Vor}^* (shown in transparent light gray).



(b) Standard Wiener Restoration using the filter in equation (1.58)



(c) Wiener Restoration for sampled images using the filter in equation (1.60)

Figure 1.17: Deconvolution of the image in figure 1.15 by two kinds of Wiener Restoration.

and then perform the same minimization as in equation (1.57) under the same statistical assumptions as before we get

$$\hat{k}(\xi) = \frac{\overline{H}(\xi)}{|H(\xi)|^2(1 + a^2(\xi) + b^2(\xi))}. \quad (1.60)$$

A similar approach has been proposed by Pratt [139] more than 20 years ago, but has since then been apparently forgotten. However, as it is shown in figure 1.17(c) this simple procedure manages to reduce much of the annoying aliasing present in the classical Wiener restoration shown in figure 1.17(b). As it can be observed, some ringing artifacts are still present, due to the zero-padding. But this can be solved by the Total Variation approach we present in the following paragraph.

Total Variation on adapted reciprocal cell. A better approach is to apply Total Variation minimization introduced by Rudin, Osher, and Fatemi [152], which is specially well suited for this kind of problems where the spectrum is known on a region which consists of several connected components, as proposed by Malgouyres [121], Malgouyres and Guichard [122]. In fact this is a good means of choosing the unknown Fourier coefficients in a way that minimizes the ringing artifacts that would otherwise be created by the gaps between two connected components.

Specifically, the most straightforward method of deconvolution and spectral extrapolation by TV minimization, would consist of minimizing $TV(u)$ under the constraint that the sample difference $\|\Delta_\Gamma \cdot ((h * u) - g)\|$ be smaller than the noise level $\|n\|$. Equivalently, we want to find u which minimizes

$$\int |\nabla u(x)| dx + \lambda \sum_{x \in \Gamma} |(h * u)(x) - g(x)|^2 \quad (1.61)$$

where λ is a properly chosen Lagrange multiplier.

However, as it has been observed in [121, 122], this procedure may lead to pixelization artifacts, especially when the spatial support of the filter h is small with respect to the sampling rate. In the extreme case where h is a delta function, the solution u of (1.61) may also become a series of delta functions at certain areas, which is certainly not what we want. Even in more realistic cases, such as the filter h corresponding to a CCD array with perfect optics, the pixelization artifacts of (1.61) may be very strong.

To avoid this problem, Malgouyres and Guichard propose to rewrite the data-fitting term in the Fourier domain as follows

$$u = \arg \inf_u \int |\nabla u(x)| dx + \lambda \int_{D_{\text{Vor}}^*} |H\hat{u} - \hat{g}|^2 d\xi. \quad (1.62)$$

The modification we propose is to integrate the data-fitting term over the $D_{\text{R-opt}}^*$ (after periodic extension of \hat{g}) instead of over D_{Vor}^*

$$u = \arg \inf_u \int |\nabla u(x)| dx + \lambda \int_{D_{\text{R-opt}}^*} |H\hat{u} - \hat{g}|^2 d\xi. \quad (1.63)$$

The benefit of doing so becomes evident when we compare the solutions of equations (1.62) and (1.63), as shown in figures 1.15(c) and 1.16(c) respectively.

Note how the proposed method produces a correct deconvolution, and even allows to extrapolate the spectrum, whereas the solution of equation (1.62) produces a catastrophic result in this case. The reason is that the high-energy second harmonic of the roof texture which falls outside of D_{Vor}^* , say at $\xi \notin D_{\text{Vor}}^*$ is translated (due to sampling and aliasing) to a region very near one of the zero-crossings of H , say at $\xi + \gamma^* \in D_{\text{Vor}}^*$. This means that $\hat{g}(\xi + \gamma^*) \approx H(\xi)\hat{f}(\xi)$ since $H(\xi + \gamma^*) \approx 0$. Hence equation (1.62), which doesn't know about aliasing, tries to deconvolve the high-energy aliased coefficient $\hat{g}(\xi + \gamma^*)$ with $H(\xi + \gamma^*)$ which is almost zero, producing the instability that we can observe in the result. Equation (1.63) instead, deconvolves $\hat{g}(\xi + \gamma^*)$ with $H(\xi)$, thus producing the correct result, closer to $\hat{f}(\xi)$.

This kind of error is not very common in more traditional systems where the transfer function is quite isotropic or the noise level too large with respect to aliasing, and hence (according to the last observation of the previous section) the optimal reciprocal cell is included in it. Furthermore this effect can be reduced by adding a threshold, which avoids to deconvolve coefficients where H is below the threshold. But this is not a very efficient solution, because it consists of treating aliasing as if it was white noise, and for highly anisotropic transfer functions with many zero-crossings, it may force us to use a quite high threshold, thus forcing us to unnecessarily throw away a lot of information.

Sometimes even in some more traditional systems addressing aliasing may be important. Figure 1.18 shows what happens when we try to deconvolve an image that has been obtained by the ccd3 system, a CCD array with good optics and a signal-to-noise ratio of 48 dB, which moved during acquisition a distance of 2.35 pixels, at constant speed, at a direction slightly deviated from the horizontal. We assume that we try to deconvolve the whole MTF including the sensor, the optics and the motion blur. It can be observed that some very annoying aliasing remains if we apply the naïve restoration, which is removed if we apply the aliasing-aware restoration.

Same comment for figure 1.19 which illustrates the same effect on a Supermode system at a signal-to-noise ratio of 48 dB, which is higher than the actual SNR of SPOT5.

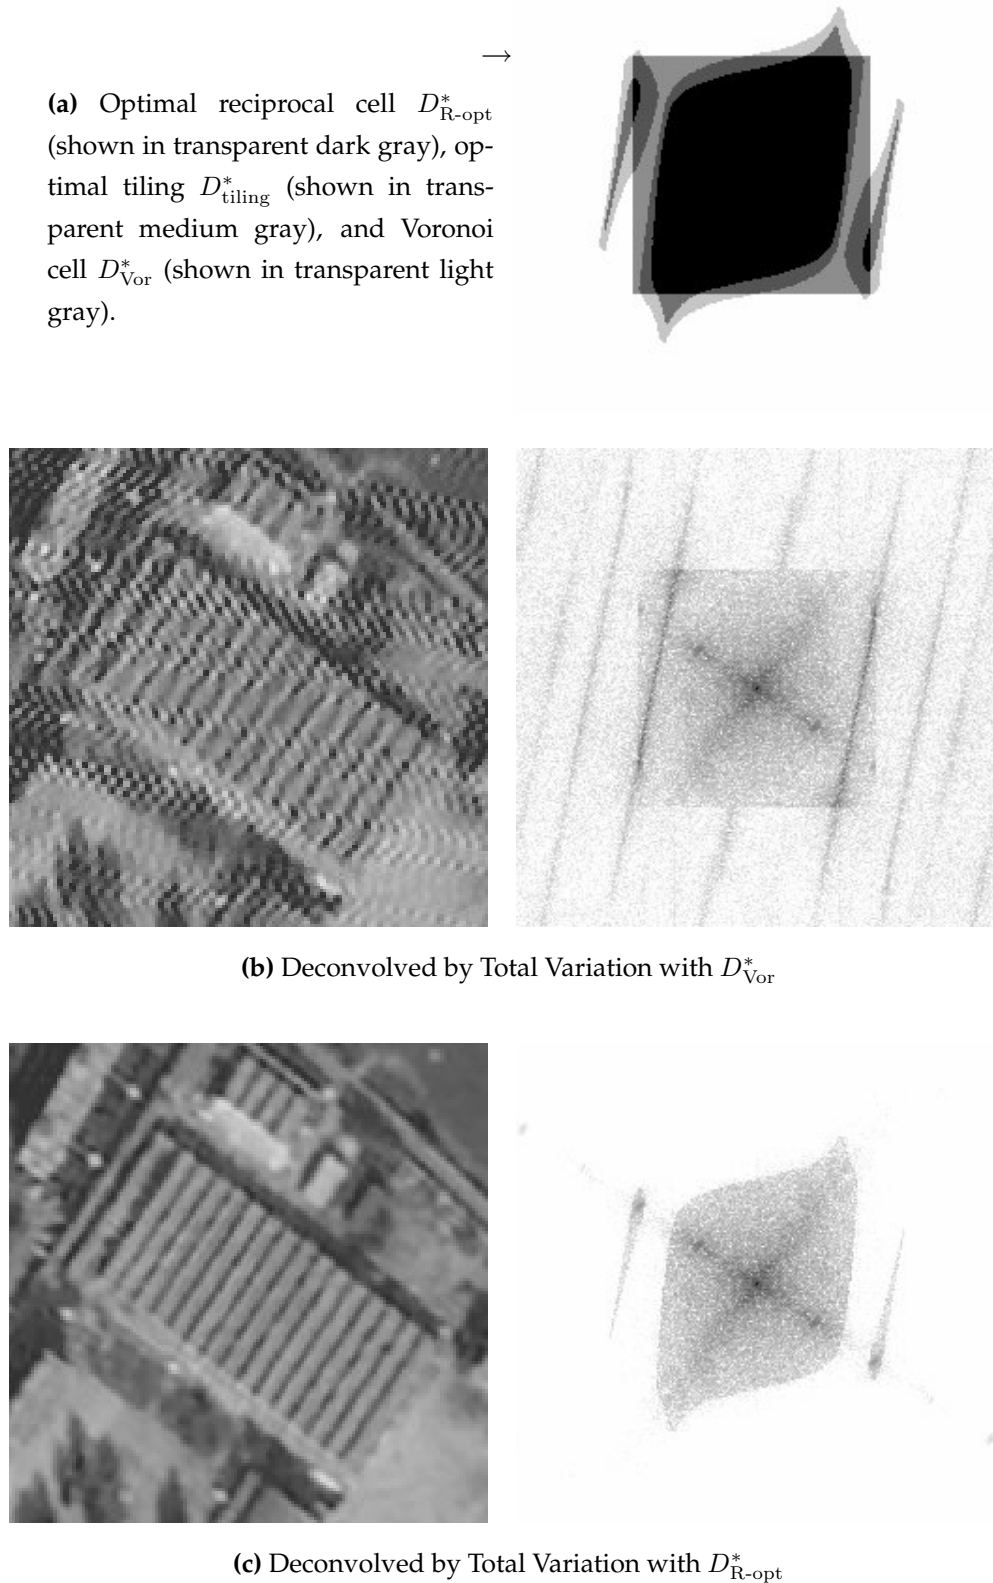





Figure 1.18: Deconvolution of a rotated and scaled version of figure 1.15(a) after sampling with the ccd3 system with a SNR of 48 dB, by Total Variation minimization with the data-fitting term integrated first over the Voronoi cell D_{Vor}^* , and then over the optimal reciprocal cell D_{R-opt}^* . Observe how the aliasing errors are reduced by placing aliasing-dominated regions at their correct locations in the Fourier spectrum.

→

(a) Optimal reciprocal cell D_{R-opt}^* (shown in transparent dark gray), optimal tiling D_{tiling}^* (shown in transparent medium gray), and Voronoi cell D_{Vor}^* (shown in transparent light gray).

(b) Deconvolved by Total Variation with D_{Vor}^*



(c) Deconvolved by Total Variation with D_{R-opt}^*

Figure 1.19: Deconvolution of a rotated and scaled version of figure 1.15(a) after sampling with the Supermode system with a SNR of 48 dB, by Total Variation minimization with the data-fitting term integrated first over the Voronoi cell D_{Vor}^* , and then over the optimal reciprocal cell D_{R-opt}^* . Observe how the aliasing errors are reduced by placing aliasing-dominated regions at their correct locations in the Fourier spectrum.

With the actual SNR aliasing is largely below the noise level, and there is not a big difference between both procedures. Note that in both cases, in order to obtain a clearly visible result we rotated and scaled the original image, in such a way that a high energy harmonic falls in the right place.


Importance of aliasing vs. noise in image restoration. Finally we show an example which illustrates the different roles played by the thresholds θ_{alias} and θ_{noise} used in the resolution measure. Figure 1.20 shows a synthetic image which was chosen to better illustrate our results. It is a radial chirp image centered at x_0 , where the local frequency at a point x is around $x - x_0$. Thus, the image and its spectrum have a very similar shape. We simulated a sampling of this image with the Sq2 system, and we applied the Equation (1.63) with two different values of R-opt. In the first case we set $\theta_{\text{alias}} = \theta_{\text{noise}} = 1$ as before, and in the second case we set $\theta_{\text{alias}} = 0.3$ whereas $\theta_{\text{noise}} = 3$, *i.e.* we tolerate higher distortions from noise than from aliasing. The result, as anticipated at the end of section 4.3, confirms that this is a sensible choice, because we perceive better distortions due to image-correlated aliasing than distortions due to image-independent noise. Another reason may be that TV too may do a better job at filtering white noise than aliasing errors.

7 Discussion and future work

In this chapter we propose a new way of measuring the effective resolution of an image acquisition system, which takes into account the noise level, and the geometric distribution of sensors and their transfer functions. This definition is motivated by modern imaging systems, which may consist of hexagonal or square grids of sensors, and may show a better resolution in certain preferred directions.

Our experiments suggest that the proposed measure can not only be used to compare the resolution of different imaging systems, but also to avoid aliasing artifacts that may cause serious problems in restoration. When this restoration is done in conjunction with total variation minimization in order to avoid ringing by fill-in the spectrum between connected components of the optimal reciprocal cell, we actually extrapolate the spectrum in a reasonable manner with a visual effect of sharpness and increased resolution.

However this spectrum extrapolation should not be confused with other “super-resolution” methods like in [78, 144], where a very detailed knowledge of image



(a) Original image and its Fourier spectrum



(b) Deconvolved by Total Variation with D_{R-opt}^* , $\theta_{alias} = \theta_{noise} = 1$ (c) Deconvolved by Total Variation with D_{R-opt}^* , $\theta_{alias} = 0.3$ and $\theta_{noise} = 3$

Figure 1.20: Deconvolution by Total Variation minimization with the data-fitting term integrated over the optimal reciprocal cell D_{R-opt}^* , for two different values of the thresholds. Observe that to completely eliminate visible aliasing we need an aliasing threshold θ_{alias} smaller than one. On the other hand, deconvolved noise is less visible or better dealt with by TV and we can tolerate higher values of the threshold θ_{noise} without any serious visible artifacts.

statistics is used to produce plausible interpolations. In our work we used a minimal knowledge of image statistics, namely the decay rate of Fourier coefficients, but we used a very detailed knowledge of the image acquisition device, its geometry, transfer function and noise. On the other hand this kind of knowledge is usually not used in super-resolution algorithms, in any case not the aliasing part.

Thus there is a potential to integrate both approaches to obtain better super-resolution methods which exploit both detailed image statistics and image acquisition models. The mutual information density measure and the corresponding reciprocal cell provide a first indication on how these approaches could be integrated. We calculated this measure to include up to second order statistics (through the relative covariance k between pairs of Fourier coefficients), but this model can be refined as much as we want. Nevertheless, image statistics are more relevant when expressed in more local bases, such as wavelets or wavelet packets bases as in [99, 144], and this would oblige us to approximate the effect of the transfer function and aliasing in such bases. Such an integration of both approaches to super-resolution will be the subject of future research.

Appendix

A Proof of Shannon's sampling theorem

In this section we give a more rigorous proof of Shannon's sampling theorem. The main elements of the theory and the proof were obtained from [36, 80, 135]. We rewrite the proofs here to unify the notation, and because we state the theorem under slightly different hypothesis which are more useful for the purposes of this chapter. In [36] for instance we can find all the elements of Fourier and distribution theory in multiple dimension, but not a proof of the sampling theorem. On the other hand [80] gives a proof of the sampling theorem in L^2 but in dimension 1. Finally [135] constructs the theory of Fourier series in several dimensions and gives a proof of the sampling theorem under conditions very similar to 5.

In the next section we recall the basic elements of the theory that we will use in section A.2 to give the proof. Finally in section A.3 we discuss to what an extent the hypothesis we assumed on the continuous image are verified by real systems.

A.1 Basic results from Fourier analysis and distribution theory

The Fourier transform and its conjugate are defined initially for $f \in L^1(\mathbb{R}^d)$ as

$$\hat{f}(\xi) = \mathcal{F}(f)(\xi) := \int_{\mathbb{R}^d} e^{-i\langle x, \xi \rangle} f(x) dx \quad (1.64)$$

$$\overline{\mathcal{F}}(F)(x) := \int_{\mathbb{R}^d} e^{i\langle x, \xi \rangle} F(\xi) d\xi. \quad (1.65)$$

The Fourier transform maps L^1 to continuous functions which tend to 0 at infinity, but $\mathcal{F}(f)$ is not necessarily in L^1 . Whenever $\hat{f} \in L^1$ we have the inversion formula

$$\overline{\mathcal{F}}(\mathcal{F}(f)) = (2\pi)^d f, \text{ if } f \in L^1, \hat{f} \in L^1.$$

On the contrary the Schwartz class \mathcal{S} of functions with fast decay at infinity for all derivatives, is closed under \mathcal{F} , so:

$$f \in \mathcal{S} \implies \mathcal{F}(f) \in \mathcal{S} \wedge \overline{\mathcal{F}}(\mathcal{F}(f)) = (2\pi)^d f.$$

The Schwartz class allows us to extend the definition of the Fourier transform to more general classes. First we can extend it to the class of tempered distributions \mathcal{S}' , so that for $u \in \mathcal{S}'$, the Fourier transform is $\hat{u} = \mathcal{F}(u)$ such that

$$\begin{aligned} \forall \varphi \in \mathcal{S}, \langle \mathcal{F}(u), \varphi \rangle &= \langle u, \mathcal{F}(\varphi) \rangle \\ \langle \overline{\mathcal{F}}(u), \varphi \rangle &= \langle u, \overline{\mathcal{F}}(\varphi) \rangle \end{aligned}$$

The definition is consistent with the definition in $L^1 \subset \mathcal{S}$, and it isometrically maps \mathcal{S}' into itself, and satisfies the Fourier inversion formula, but now in the sense of tempered distributions:

$$f \in \mathcal{S}' \implies \mathcal{F}(f) \in \mathcal{S}' \wedge \quad (1.66)$$

$$\overline{\mathcal{F}}(\mathcal{F}(f)) = (2\pi)^d f. \quad (1.67)$$

The extension of the Fourier transform to \mathcal{S}' provides an extension to $L^2(\mathbb{R}^d) \subset \mathcal{S}'(\mathbb{R}^d)$. Since \mathcal{S} is complete and dense in L^2 we can show that this is the only extension of $\mathcal{F} : \mathcal{S} \rightarrow \mathcal{S}$. We can check that this extension maps L^2 to itself, that the map is continuous, that it preserves the L^2 norm up to a factor $(2\pi)^{\frac{d}{2}}$ and that it satisfies the Fourier inversion formula (in the sense of the L^2 norm):

$$f \in L^2 \implies \mathcal{F}(f) \in L^2 \wedge \quad (1.68)$$

$$\overline{\mathcal{F}}(\mathcal{F}(f)) = (2\pi)^d f \wedge \quad (1.69)$$

$$\|\mathcal{F}(f)\|_{L^2} = (2\pi)^{\frac{d}{2}} \|f\|_{L^2}$$

The Fourier transform in L^2 is closely related to the more elementary notion of Fourier series of periodic functions. Any Γ -periodic function $f \in L^2_\Gamma$ admits a Fourier series development

$$f(x) = \sum_{k \in \Gamma^*} c_k e^{i\langle k, x \rangle} \quad (1.70)$$

where equality holds in the sense of the $L^2(D)$ norm and the Fourier coefficients are unique and defined by

$$c_k = \frac{1}{S} \int_D e^{-i\langle x, k \rangle} f(x) dx$$

for any tile D of Γ , where $S = |D|$ denotes the area of any such tile. The link with the Fourier transform of distributions is given by:

$$\hat{u} = (2\pi)^d \sum_{k \in \Gamma^*} c_k(u) \delta_k$$

which is valid for any Γ -periodic function u .

Similarly, convolution is defined by the usual formula for $f, g \in L^1(\mathbb{R}^d)$

$$(f * g)(x) = \int_{\mathbb{R}^d} f(x-y)g(y)dy$$

and then extended in several ways. The main difficulty is that the decay rate at infinity of both terms need to compensate for the integral to converge. For our purposes it suffices to know that it is well defined in the following cases

$$\begin{aligned} \forall u \in \mathcal{E}', \forall \varphi \in C_0^\infty, (u * \varphi) \in C_0^\infty \text{ with} \\ \text{supp}(u * \varphi) = \text{supp}(u) + \text{supp}(\varphi) \end{aligned} \quad (1.71)$$

$$\forall u \in \mathcal{E}', \forall v \in \mathcal{S}', (u * v) \in \mathcal{S}'$$

$$\forall u \in \mathcal{O}'_C, \forall v \in \mathcal{S}', (u * v) \in \mathcal{S}'$$

$$\forall u \in \mathcal{S}, \forall v \in \mathcal{S}', (u * v) \in \mathcal{O}_M.$$

Our notation follows [36]: \mathcal{E}' denotes the compact support distributions, \mathcal{O}'_C the distributions with fast decay at infinity, and \mathcal{O}_M the functions having slow growth at infinity and an infinite number of continuous derivatives. Observe that we have the following inclusions relationships between these spaces:

$$\begin{array}{ccccccc} C_0^\infty & \longrightarrow & \mathcal{S} & \longrightarrow & \mathcal{O}_M & \longrightarrow & C^\infty \\ \downarrow & & \downarrow & & \downarrow & & \downarrow \\ \mathcal{E}' & \longrightarrow & \mathcal{O}'_C & \longrightarrow & \mathcal{S}' & \longrightarrow & \mathcal{D}' \end{array} \quad (1.72)$$

In addition $\mathcal{S} \subset L^p \subset \mathcal{S}'$ for all $p \geq 1$. In all the previous cases where the convolution is well defined, its Fourier transform exchanges with convolution in the following manner:

if u and v satisfy one of the following conditions:

$$u \in \mathcal{E}' \wedge v \in \mathcal{S}'$$

$$u \in \mathcal{O}'_C \wedge v \in \mathcal{S}'$$

$$u \in \mathcal{S} \wedge v \in \mathcal{S}'$$

$$\text{then } \mathcal{F}(u * v) = \mathcal{F}(u) \cdot \mathcal{F}(v). \quad (1.73)$$

which is also valid for $\overline{\mathcal{F}}$ instead of \mathcal{F} .

The multiplication of two distributions is not always well defined either. For our purposes it suffices to know that we can safely multiply a tempered distribution by a slowly growing regular function

$$f \in \mathcal{O}_M \wedge g \in \mathcal{S}' \implies f \cdot g \in \mathcal{S}' \quad (1.74)$$

and that the Fourier transform maps distributions with compact support into \mathcal{O}_M :

$$f \in \mathcal{E}' \implies \mathcal{F}(f) \in \mathcal{O}_M. \quad (1.75)$$

These two results allow to obtain the dual form of equation (1.73), namely:

if u and v satisfy one of the following conditions:

$$\hat{u} \in \mathcal{E}' \wedge v \in \mathcal{S}'$$

$$u \in \mathcal{O}_M \wedge \hat{u} \in \mathcal{O}'_C \wedge v \in \mathcal{S}'$$

$$u \in \mathcal{S} \wedge v \in \mathcal{S}'$$

$$\text{then } \mathcal{F}(u \cdot v) = \frac{1}{(2\pi)^d} \mathcal{F}(u) * \mathcal{F}(v). \quad (1.76)$$

To conclude we state two more results that we will need in the following section. The first is valid for $f \in \mathcal{S}'$, $f \in \mathcal{S}$ or $f \in L^2$, with the corresponding meaning of equality in each case:

$$\mathcal{F}(f(x)e^{-i\langle x, \omega \rangle}) = \tau_\omega \hat{f} \quad (1.77)$$

$$\mathcal{F}(\Delta_\Gamma) = S^* \Delta_{\Gamma^*}. \quad (1.78)$$

In the last line $S^* = |D^*|$ denotes the area of any tile D^* of Γ^* .

A.2 Shannon's sampling theorem

The formal proof given in section 2.2 is actually only valid if we impose a strong regularity and fast decay on the analog image f :

Theorem 3 (Shannon-Whittaker in \mathcal{S}). *Given a function $f \in \mathcal{S}(\mathbb{R}^d)$ (with fast decay and an infinite number of continuous derivatives), a d -dimensional sampling grid Γ and a compact reciprocal cell $D^* \subset \mathbb{R}^d$. If the following conditions are met*

(S1) (D^*, Γ^*) is a packing

(S2) $\text{supp}(\hat{f}) \subset D^*$ (i.e. f is band-limited),

then f can be completely recovered from its samples in Γ , i.e. from

$$g = \{f_\gamma\}_{\gamma \in \Gamma}. \quad (1.79)$$

In fact the reconstruction is given by the following convolution

$$f = \frac{1}{S^*} \overline{\mathcal{F}}(\mathbb{1}_{D^*}) * g \quad (1.80)$$

of the sampled image g with the generalized sinc kernel $s = \frac{1}{S^*} \overline{\mathcal{F}}(\mathbb{1}_{D^*})$.

Proof. Since $f \in \mathcal{S}$ (which is included in \mathcal{O}_M) and $\Delta_\Gamma \in \mathcal{S}'$ the product $g = \Delta_\Gamma \cdot f$ defining the sampled image g makes sense and is in \mathcal{S}' according to (1.74).

Now again, since f has compact spectral support we can apply (1.76) in the case when $u = f$ is such that $\hat{u} \in \mathcal{E}'$ and $v = \Delta_\Gamma \in \mathcal{S}'$ to obtain: $\mathcal{F}(f \cdot \Delta_\Gamma) = \frac{1}{(2\pi)^d} \mathcal{F}(f) * \mathcal{F}(\Delta_\Gamma)$. Finally applying (1.78) to compute the Fourier transform of the Dirac comb the previous equation becomes Poisson's formula:

$$\mathcal{F}(f \cdot \Delta_\Gamma) = \frac{S^*}{(2\pi)^d} \hat{f} * \Delta_{\Gamma^*},$$

where both members are tempered distributions, and equality is to be understood in this sense.

Multiplying both sides by $\mathbb{1}_{D^*} \in \mathcal{O}_M$ we still get the following equality in \mathcal{S}'

$$\mathcal{F}(f \cdot \Delta_\Gamma) \cdot \mathbb{1}_{D^*} = \frac{S^*}{(2\pi)^d} (\hat{f} * \Delta_{\Gamma^*}) \cdot \mathbb{1}_{D^*} = \frac{S^*}{(2\pi)^d} \hat{f}. \quad (1.81)$$

Regarding the last equality observe that (S1) and (S2) ensure that there is no aliasing, in other words, that $\text{supp}(\hat{f} * \delta_{\gamma^*}) \cap D^* = \emptyset$ for any $\gamma^* \in \Gamma^* \setminus \{0\}$. Thus $(\hat{f} * \Delta_{\Gamma^*}) \cdot \mathbb{1}_{D^*} = \sum_{\gamma^* \in \Gamma^*} (\hat{f} * \delta_{\gamma^*}) \cdot \mathbb{1}_{D^*}$ actually has only one non-zero term, namely when $\gamma^* = 0$.

Finally we only have to compute the inverse Fourier transform on both sides of (1.81). According to the Fourier inversion formula (1.67) the right-hand side gives $S^* f$, and to compute the left-hand side we apply once again (1.76) to obtain

$$\frac{1}{(2\pi)^d} \overline{\mathcal{F}}(\mathcal{F}(g)) * \overline{\mathcal{F}}(\mathbb{1}_{D^*}) = S^* f.$$

Observe that we can apply (1.76) for $u = \mathcal{F}(g)$ and $v = \mathbb{1}_{D^*} \in \mathcal{S}'$, because $g \in \mathcal{O}'_C$. In fact, g is clearly a distribution and it has a fast decay because so does $f \in \mathcal{S}$.

To complete the proof we only need to apply once more the Fourier inversion formula to obtain $\overline{\mathcal{F}}(\mathcal{F}(g)) = (2\pi)^d g$ and rearrange the terms in the last equation. \square

Note that this proof shows the interpolation equation (1.80) only in the sense of \mathcal{S}' (tempered distributions). We shall see later that it is also valid pointwise with uniform convergence and with less restrictive hypothesis on the decay rate of the continuous signal f at infinity. Part of the proof, however, will require results from Fourier series in L^2 .

We can still give a proof which is completely based on distribution arguments if we trade some decay speed in f for a slightly over-critical sampling which allows for interpolation kernels s with a more regular spectrum. This form of Shannon's sampling theorem is also interesting for the numerical implementation, since the more regular form of \hat{s} means a faster decay rate for s and hence faster approximate interpolation.

Theorem 4 (Shannon-Whittaker in \mathcal{O}_M). *Given a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, a d -dimensional sampling grid Γ and a compact reciprocal cell $D^* \subset \mathbb{R}^d$. If*

- (S1) (D^*, Γ^*) is a packing
- (S2') $\text{supp}(\hat{f}) \subset \overset{\circ}{D^*}$ (i.e. f is band-limited, and D^* is slightly over-critical),

then f can be completely recovered from its samples in Γ , i.e. from

$$g = \Delta_\Gamma \cdot f. \tag{1.82}$$

In fact the reconstruction is given by the following convolution

$$f = \frac{1}{S^*} \overline{\mathcal{F}}(\varphi) * g \tag{1.83}$$

of the sampled image g with any kernel $s = \frac{1}{S^*} \overline{\mathcal{F}}(\varphi)$ satisfying the following conditions:

- (S3) $\varphi \in C_0^\infty$, $\text{supp}(\varphi) \subseteq D^*$, $\varphi|_{\text{supp}(f)} \equiv 1$.

Such a kernel always exists but is not unique.

Proof. Since f has compact spectral support, according to equation (1.75), it has an infinite number of continuous derivatives and grows slowly, i.e. $f \in \mathcal{O}_M$. Since in addition $\Delta_\Gamma \in \mathcal{S}'$ the product $g = \Delta_\Gamma \cdot f$ defining the sampled image g makes sense and is in \mathcal{S}' according to (1.74).

Then we derive Poisson's formula exactly as before:

$$\mathcal{F}(f \cdot \Delta_\Gamma) = \frac{S^*}{(2\pi)^d} \hat{f} * \Delta_{\Gamma^*}.$$

Now let's construct φ satisfying (S3). Let $K = \text{supp}(\hat{f})$ denote the compact spectral support of f . From (S2) there is a positive distance $d(\partial D^*, K) = \varepsilon > 0$ between the border of D^* and K . Let $\psi \in C_0^\infty$ be a test function with support in a zero-centered ball of radius $\varepsilon/2$, and let K_ε be the dilation of K by the same ball. Then it is easy to see from (1.71) that $\varphi = \psi * \mathbb{1}_{K_\varepsilon}$ satisfies all conditions in (S3).

Multiplying both sides of the Poisson formula by $\varphi \in C_0^\infty$ we still get the following equality in \mathcal{S}'

$$\mathcal{F}(f \cdot \Delta_\Gamma) \cdot \varphi = \frac{S^*}{(2\pi)^d} (\hat{f} * \Delta_{\Gamma^*}) \cdot \varphi = \frac{S^*}{(2\pi)^d} \hat{f}. \quad (1.84)$$

The last equality is shown in exactly the same manner as before. The only non-zero term in the convolution is when $\gamma^* = 0$ because $\text{supp}(\varphi) \subseteq D^*$, and this term is $(\hat{f} * \delta_0) \cdot \varphi = \hat{f}$ because $\varphi|_{\text{supp}(\hat{f})} \equiv 1$.

Finally we only have to compute the inverse Fourier transform on both sides of (1.84). According to the Fourier inversion formula (1.67) the right-hand side gives $S^* f$, and to compute the left-hand side we apply once again (1.76) to obtain

$$\frac{1}{(2\pi)^d} \overline{\mathcal{F}}(\mathcal{F}(g)) * \overline{\mathcal{F}}(\varphi) = S^* f.$$

Observe that this time we can apply (1.76) for $v = \mathcal{F}(g)$ and $u = \varphi$, but the image and the kernel play inverse roles as in the previous case. As for the sampled image g the best thing we can say is that it grows slowly, but it is not in \mathcal{O}_M since it is a distribution. So $g \in \mathcal{S}'$ is a tempered distribution and so is $v = \mathcal{F}(g) \in \mathcal{S}'$ according to (1.66). Thus the necessary decay rate for the formula to hold is provided by $u = \varphi \in C_0^\infty \subset \mathcal{S}$.

To complete the proof we only need to apply once more the Fourier inversion formula to obtain $\overline{\mathcal{F}}(\mathcal{F}(g)) = (2\pi)^d g$ and rearrange the terms in the last equation. \square

Finally we provide an optimal formulation we know of Shannon's sampling theorem in the case of critical sampling. It is optimal in the sense that it holds with the weakest hypothesis, and that it provides the strongest results.

Theorem 5 (Shannon-Whittaker in L^2). Given a function $f \in L^2(\mathbb{R}^d)$, a d -dimensional sampling grid Γ and a compact reciprocal cell $D^* \subset \mathbb{R}^d$. If the following conditions are met

(S1) (D^*, Γ^*) is a packing

(S2) $\text{supp}(\hat{f}) \subset D^*$ (i.e. f is band-limited),

then f can be completely recovered from its samples in Γ , i.e. from

$$g = \{f(\gamma)\}_{\gamma \in \Gamma}. \quad (1.85)$$

In fact, $g \in l^2(\Gamma)$, and the reconstruction is given by the following convolution

$$f(x) = \sum_{\gamma \in \Gamma} g_\gamma s(x - \gamma) \quad (1.86)$$

of the sampled image g with the generalized sinc kernel $s = \frac{1}{S^*} \overline{\mathcal{F}}(\mathbb{1}_{D^*})$. This equality holds in general only in the sense of the L^2 norm. If in addition $g \in l^1(\Gamma)$, then there is pointwise equality and uniform convergence.

Proof. Since f has compact spectral support, $\hat{f} \in \mathcal{E}'$ and we deduce Poisson's formula in the same manner as in theorem 3

$$\mathcal{F}(f \cdot \Delta_\Gamma) = \frac{S^*}{(2\pi)^d} \hat{f} * \Delta_{\Gamma^*},$$

where both members are tempered distributions, and equality is to be understood in this sense.

Now observe that the right-hand side of Poisson's formula $F = \frac{S^*}{(2\pi)^d} \hat{f} * \Delta_{\Gamma^*}$ is a Γ^* -periodic function in $L^2(R)$ for any tile R of Γ^* (this is due to the fact that $f \in L^2(\mathbb{R}^d)$ and hence also $\hat{f} \in L^2(\mathbb{R}^d)$, according to (1.68)). Therefore F admits a Fourier series development:

$$F(\xi) = \sum_{\gamma \in \Gamma} c_\gamma e^{i\langle \xi, \gamma \rangle} \quad (1.87)$$

where $\{c_{-\gamma}\}_\Gamma$ is in $l^2(\Gamma)$, and the equality is valid in the $L^2(R)$ for any tile R of Γ^* .

On the other hand the left-hand side of the Poisson formula can be written:

$$F(\xi) = \sum_{\gamma \in \Gamma} f(\gamma) e^{-i\langle \gamma, \xi \rangle}. \quad (1.88)$$

So the uniqueness of the Fourier series development (1.70) means that $c_{-\gamma} = f(\gamma) = g_\gamma$, and we conclude that the sampled image $g \in l^2(\Gamma)$ is square-summable.

Multiplying F by $\mathbb{1}_{D^*} \in L^2$ we still get the following equality in $L^2(\mathbb{R}^d)$

$$\sum_{\gamma \in \Gamma} f(\gamma) \mathbb{1}_{D^*}(\xi) e^{-i\langle \gamma, \xi \rangle} = F(\xi) \mathbb{1}_{D^*}(\xi) = \frac{S^*}{(2\pi)^d} ((\hat{f} * \Delta_{\Gamma^*}) \cdot \mathbb{1}_{D^*})(\xi) = \frac{S^*}{(2\pi)^d} \hat{f}(\xi). \quad (1.89)$$

To show the last equality observe as in theorem 3 that (S1) and (S2) ensure that there is no aliasing.

Finally we only have to compute the inverse Fourier transform on both sides of (1.81). According to the Fourier inversion formula (1.69) the right-hand side gives $S^* f$, and to compute the left-hand side we use the continuity of $\overline{\mathcal{F}}$ (1.68) to obtain

$$\overline{\mathcal{F}}\left(\sum_{\gamma \in \Gamma} f(\gamma) \mathbb{1}_{D^*}(\xi) e^{-i\langle \gamma, \xi \rangle}\right) = \sum_{\gamma \in \Gamma} f(\gamma) \overline{\mathcal{F}}(\mathbb{1}_{D^*}(\xi) e^{-i\langle \gamma, \xi \rangle}).$$

Using equation (1.77) we substitute $\overline{\mathcal{F}}(\mathbb{1}_{D^*}(\xi) e^{-i\langle \gamma, \xi \rangle}) = \overline{\mathcal{F}}(\mathbb{1}_{D^*})(x - \gamma)$, to obtain:

$$f(x) = \frac{1}{S^*} \sum_{\gamma \in \Gamma} g_\gamma \overline{\mathcal{F}}(\mathbb{1}_{D^*})(x - \gamma)$$

which is equivalent to (1.86), and the convergence of the series towards f holds in the sense of the $L^2(\mathbb{R}^d)$ norm. If in addition $g \in l^1(\Gamma)$ (i.e. $\sum_\gamma |g_\gamma| < +\infty$) then, since $s = \frac{1}{S^*} \overline{\mathcal{F}}(\mathbb{1}_{D^*})$ is bounded¹⁰, we also have $\sum_{|\gamma| < n} |g_\gamma s(x - \gamma)| \leq \sum_{|\gamma| < n} |g_\gamma| = c_n \rightarrow 0$ (where the sequence c_n is independent of x), Therefore the series converges uniformly to a function h which is continuous over \mathbb{R}^d . So, the series converges also to h in $L^2(J)$ for any bounded interval J . We conclude that $h = f$ almost everywhere over \mathbb{R}^d , which means that $f(x) = h(x)$ for all $x \in \mathbb{R}^d$, since f and h are continuous. \square

A.3 Suitability of the hypotheses in real systems

As we explained in section 3 in real image acquisition systems the continuous image f to be sampled can be modeled as a convolution:

$$f = h * O \quad (1.90)$$

where O represents the landscape, i.e. the photons that meet the imaging system, and the kernel h depends on the filtering performed by the optics of the acquisition device, the sensors and the movement of the whole system. Thus O can be modeled as a positive Radon measure with compact support, and hence $O \in \mathcal{E}'$. As for the kernel h , despite the approximate models introduced in section 3, it is commonly accepted that

¹⁰ In fact $|s(x)| \leq \frac{1}{S^*} \int_{\mathbb{R}^d} |\mathbb{1}_{D^*}(\xi) e^{-i\langle \xi, x \rangle}| d\xi \leq \frac{1}{S^*} \int_{D^*} 1 d\xi = \frac{|D^*|}{S^*} \leq 1$, where the last inequality holds because D^* is a sub-tile of Γ^* .

it has compact spectral support. However this support is usually much larger than its essential support (beyond the noise level), so from an engineering point of view this remark is not very useful. But it is useful for checking the validity of the theorem for real images.

Now since $\hat{h} \in \mathcal{E}'$ we conclude that $\hat{f} = \hat{h} \cdot \hat{O}$ is also in \mathcal{E}' according to (1.73), and $f \in \mathcal{O}_M$. So this is enough to satisfy the conditions of theorem 4 with over-critical sampling.

To meet the hypothesis of the two other versions we shall need to assume some more regularity on the kernel h . For theorem 4 we shall need to assume that the transfer function $\hat{h} \in C_0^\infty$ is a test function. Then according to (1.73) we have $\hat{f} = \hat{h} \cdot \hat{O}$ with $\hat{O} \in \mathcal{O}_M$. Therefore f has both compact support and an infinite number of derivatives, so $\hat{f} \in C_0^\infty \subset \mathcal{S}$, and $f \in \mathcal{S}$ according to (A.1), so f meets the hypothesis in theorem 4.

Finally, for theorem 5 it suffices to assume that \hat{h} is continuous (and in addition of compact support as before). So it is in $L^1 \cap L^2$, and $\hat{h} \in \mathcal{O}_M \cap L^2$ according to equations (1.68) and (1.75). We can conclude that $f = h * O$ is C^∞ that it tends to zero at infinity and $f \in L^2$.

B Proof of the lower bound on mutual information density

The proof of theorem 2 is a simple application of the following results from information theory, which can be found for instance in [51].

Lemma 1. *Given any three random variables X_1, X_2, X_3*

$$H(X_1 + X_2 + X_3) \geq H(Y_1 + Y_2 + Y_3) \text{ where } Y_i \text{ are indep. normal} \quad (1.91)$$

$$\text{with } H(X_i) = H(Y_i)$$

$$\text{and } = \text{ iff } X_i \text{ are indep. normal.}$$

Lemma 2. *Given any pair of random variables X_1, X_2 :*

$$H(X_1 + X_2 | X_1) = H(X_2 | X_1) \leq H(X_2) \quad (1.92)$$

$$\text{with } = \text{ iff } X_1, X_2 \text{ indep.}$$

Lemma 3. *The entropy of any normally distributed random vector \mathbf{X} with zero mean and covariance matrix \mathbf{K} is*

$$H(\mathbf{X}) = \frac{1}{2} \log_2 (2e\pi |\mathbf{K}|) \quad (1.93)$$

where $|\mathbf{K}|$ denotes the determinant of \mathbf{K} .

In particular, for a zero-mean, normally distributed random variable X the entropy is:

$$H(X) = \frac{1}{2} \log_2 (2e\pi \text{Var}\{X\}). \quad (1.94)$$

Lemma 4. Given any normally distributed random vector \mathbf{X} with zero mean and covariance matrix \mathbf{K} , the sum $\sum_i X_i$ is also zero-mean, normally distributed, with variance:

$$\text{Var}\left\{\sum_i X_i\right\} = \mathbf{1}^T \mathbf{K} \mathbf{1} \quad (1.95)$$

where $\mathbf{1}$ is a column vector of ones having the same size as \mathbf{X} .

In particular, given two independent, zero-mean, normally distributed random variables Y_1, Y_2 , the sum $Y_1 + Y_2$ is also zero-mean normally distributed with $\text{Var}\{Y_1 + Y_2\} = \text{Var}\{Y_1\} + \text{Var}\{Y_2\}$.

Proof. (of theorem 2)

The mutual information is defined in terms of entropy H as follows:

$$\begin{aligned} I(\hat{g}(\xi), G(\xi)) &= H(G(\xi) + G_{\text{alias}}(\xi) + \hat{n}(\xi)) - \\ &\quad H(G(\xi) + G_{\text{alias}}(\xi) + \hat{n}(\xi) | G(\xi)) \\ &= H(G(\xi) + G_{\text{alias}}(\xi) + \hat{n}(\xi)) - \\ &\quad H(G_{\text{alias}}(\xi) + \hat{n}(\xi) | G(\xi)). \end{aligned} \quad (1.96)$$

If we apply lemmas 1 and 2 to each term of the previous equation with $X_1 = G(\xi)$, $X_2 = G_{\text{alias}}(\xi)$ and $X_3 = \hat{n}(\xi)$ we obtain the following lower bound:

$$I(\hat{g}(\xi), G(\xi)) \geq H(Y_1 + Y_2 + Y_3) - H(G_{\text{alias}}(\xi) + \hat{n}(\xi)) \quad (1.97)$$

with Y_i independent and normally distributed. Hence we can apply lemmas 4 and 3 to obtain:

$$\begin{aligned} H(Y_1 + Y_2 + Y_3) &= \frac{1}{2} \log_2 (2\pi e (\text{Var}\{Y_1 + Y_2 + Y_3\})) \\ &= \frac{1}{2} \log_2 (2\pi e (\text{Var}\{Y_1\} + \text{Var}\{Y_2\} + \text{Var}\{Y_3\})) \\ &= \frac{1}{2} \log_2 (2^{2H(Y_1)} + 2^{2H(Y_2)} + 2^{2H(Y_3)}) \\ &= \frac{1}{2} \log_2 (2^{2H(G(\xi))} + 2^{2H(G_{\text{alias}}(\xi))} + 2^{2H(\hat{n}(\xi))}) \end{aligned} \quad (1.98)$$

where in the last equation we used the fact that Y_i in lemma 2 are such that $H(Y_i) = H(X_i)$. Finally if we combine this last expression with $H(G_{\text{alias}}(\xi) + \hat{n}(\xi)) = \frac{1}{2} \log_2 (2^{2H(G_{\text{alias}}(\xi) + \hat{n}(\xi))})$,

we obtain equation (1.39). Finally, a similar but simpler application of lemmas 4 and 3, shows that in the independent Gaussian case we actually have

$$2^{2\mathbf{H}(G_{\text{alias}}(\xi) + \hat{n}(\xi))} = 2^{2\mathbf{H}(G_{\text{alias}}(\xi))} + 2^{2\mathbf{H}(\hat{n}(\xi))}$$

so that the lower-bound in equation (1.39) becomes

$$\frac{1}{2} \log_2 \left(1 + \frac{2^{2\mathbf{H}(G(\xi))}}{2^{2\mathbf{H}(G_{\text{alias}}(\xi))} + 2^{2\mathbf{H}(\hat{n}(\xi))}} \right).$$

Then equation (1.40) follows directly by applying lemma 3 in the scalar case.

To prove the third part of the theorem we go back to equation (1.96). To develop the first term, in that equation, observe that the covariance matrix of $(G(\xi), G_{\text{alias}}(\xi), \hat{n}(\xi))$ consists of a 2 by 2 diagonal block (corresponding to $(G(\xi), G_{\text{alias}}(\xi))$), and a 1 by 1 diagonal block (corresponding to $\hat{n}(\xi)$). Hence:

$$\begin{aligned} \mathbf{1}^T \mathbf{K} \mathbf{1} &= \text{Var}\{G(\xi)\} + \text{Var}\{G_{\text{alias}}(\xi)\} \\ &\quad + 2\text{Cov}\{G(\xi), G_{\text{alias}}(\xi)\} + \text{Var}\{\hat{n}(\xi)\} \end{aligned} \quad (1.99)$$

and we can develop the first term as follows:

$$\begin{aligned} \mathbf{H}(G(\xi) + G_{\text{alias}}(\xi) + \hat{n}(\xi)) &= \frac{1}{2} \log(2e\pi(\text{Var}\{G(\xi)\} + \text{Var}\{G_{\text{alias}}(\xi)\} \\ &\quad + 2\text{Cov}\{G(\xi), G_{\text{alias}}(\xi)\} + \text{Var}\{\hat{n}(\xi)\})) \end{aligned} \quad (1.100)$$

To develop the second term in equation (1.96) we use the definition of conditional entropy, and then lemma 3, to obtain:

$$\begin{aligned} \mathbf{H}(G_{\text{alias}}(\xi) + \hat{n}(\xi) | G(\xi)) &= \mathbf{H}(G(\xi), G_{\text{alias}}(\xi) + \hat{n}(\xi)) - \mathbf{H}(G(\xi)) \\ &= \frac{1}{2} \log(2e\pi(\text{Var}\{G(\xi)\}\text{Var}\{G_{\text{alias}}(\xi) + \hat{n}(\xi)\} \\ &\quad - \text{Cov}\{G(\xi), G_{\text{alias}}(\xi) + \hat{n}(\xi)\}^2)) \\ &= \frac{1}{2} \log(2e\pi(\text{Var}\{G(\xi)\}(\text{Var}\{G_{\text{alias}}(\xi)\} + \text{Var}\{\hat{n}(\xi)\}) \\ &\quad - \text{Cov}\{G(\xi), G_{\text{alias}}(\xi)\}^2)). \end{aligned} \quad (1.101)$$

In the last line we applied lemma 4 to the independent Gaussian variables $G_{\text{alias}}(\xi)$, $\hat{n}(\xi)$ together with the fact that $\text{Cov}\{G(\xi), G_{\text{alias}}(\xi) + \hat{n}(\xi)\} = \text{Cov}\{G(\xi), G_{\text{alias}}(\xi)\}$ due to the independence of $G(\xi)$ and $\hat{n}(\xi)$.

Finally we combine the two terms in equation (1.96) using the previous dimensionless variables to obtain equation (1.41) which is always well defined, since $k(\xi) \in [0, 1]$, because $(\text{Var}\{G(\xi)\}\text{Var}\{G_{\text{alias}}(\xi)\} - \text{Cov}\{G(\xi), G_{\text{alias}}(\xi)\}^2)$ must be non-negative, since it is the determinant of the (positive semi-definite) covariance matrix of $(G(\xi), G_{\text{alias}}(\xi))$.

□

C Sufficient conditions for admissibility

Definition 6 for an admissible alias selector is quite general but a bit technical. Here we present a set of sufficient conditions that may be easier to verify. In particular they can be easily checked for the function u we apply it to, namely $u(\xi) = \frac{1}{2} \log_2(1 + \frac{1}{a(\xi)^2 + b(\xi)^2})$ when a and b are the relative aliasing and noise measures for a given transfer function and image and noise models.

Lemma 5. *Let Γ^* be a regular grid in \mathbb{R}^d . For any function $u : \mathbb{R}^d \rightarrow \mathbb{R}$, satisfying the following properties*

(i) $u(x) \geq 0, \forall x \in \mathbb{R}^d$

(ii) $u(x) \rightarrow 0, \text{ when } |x| \rightarrow \infty$

(iii) $(\Gamma^*, \text{supp}(u))$ is a covering

(iv) The periodicity set P_{u,Γ^*} of u and Γ^* defined below has zero measure

$$P_{u,\Gamma^*} = \{x \in \mathbb{R}^d : \exists \gamma \in \Gamma^* \setminus \{0\}, u(x) = u(x + \gamma) > 0\}$$

(iv) the associated optimal tiling R_{\max} is bounded

then u is an admissible alias selector.

Proof. Recall from the definition of admissible alias selector that the anomaly set is defined as

$$A_{u,\Gamma^*} = \{\xi \in \mathbb{R}^d : \text{card}(G_{\xi,u,\Gamma^*}) \neq 1\}$$

where $G_{\xi,u,\Gamma^*} = \left\{ \gamma_0 \in \Gamma^* : u(\xi + \gamma_0) = \sup_{\gamma \in \Gamma^*} u(\xi + \gamma) \right\}$.

For each point x consider the set G_{x,u,Γ^*} (that we abbreviate in the sequel by G_x), and let's characterize the anomaly set by the following case discussion.

Case 1: If G_x consists of a single element γ_x , then x is not in the anomaly set A_{u,Γ^*} .

Case 2: We cannot have $G_x = \emptyset$ because this would mean that there is a sequence $\{\gamma_n\}$ of elements in Γ^* , such that $\lim u(x + \gamma_n) = u_x$, and the limit u_x is strictly larger than any $u(x + \gamma_n)$. Since γ takes values in Γ^* this means that $\{\gamma_n\}$ is unbounded, so according to (ii) we have $\lim u(x + \gamma_n) = u_x = 0$. Since (i) imposes u to be non-negative this means that $u(x + \gamma_n) = u_x = 0$ for all n , which gives an infinite number of elements

in G_x , contradicting our initial supposition that $G_x = \phi$.

Case 3: Similarly, if $G_x = \{\gamma_n : n \in \mathbb{N}\}$ has an infinite number of different elements we conclude that $u_x = 0$. Let's call the set of such points zero-set of u and Γ^* :

$$Z_{u,\Gamma^*} = \{x : u_x = 0\}.$$

Clearly any point $y \in \bigcup_{\gamma \in \Gamma^*} (\text{supp}(u) + \gamma)$ is excluded from the zero-set, i.e. $y \notin Z_{u,\Gamma^*}$. Therefore, (iii) implies that the zero-set Z_{u,Γ^*} has zero measure.

Case 4: Finally if G_x contains a finite number of elements but more than one, then x belongs to the periodicity set of u and Γ^* .

We conclude that the anomaly set is the union of the zero set and the periodicity set, both of zero measure. So the anomaly set $A_{u,\Gamma^*} = P_{u,\Gamma^*} \cup Z_{u,\Gamma^*}$ also has zero measure. \square

More generally, observe that in proposition 1 the anomaly set can as well be divided into two sets: A^+ consisting of all points x such that G_x has more than one element, and A^0 consisting of all points x such that G_x is empty. Observe that we could tolerate that A^+ has a positive measure, but not that A^0 has a positive measure. In the latter case there is no maximal tiling, whereas in the former case we have several possible maximal tilings. Indeed for any $x \in A^+$, to build a maximal tiling R_{\max} , we can substitute x by $(x + \gamma)$ for any $\gamma \in G_x$, without changing the value of the integrand in $\int_{R_{\max}} v(u(x))dx$.

Chapter 2

Irregular Sampling and Reconstruction Algorithms *

Abstract

Satellite images are sampled on a slightly perturbed grid, due to micro-vibrations of the instrument during capture. This perturbation can be estimated with high accuracy, but it must be also corrected in the images for certain stereo and multi-spectral applications. In this chapter we show to what an extent future satellites being developed at CNES satisfy the conditions required by irregular sampling theory to make the problem of resampling on a regular grid well posed, whereas most current imaging systems do not allow for such a well posed reconstruction due to aliasing. Then we discuss the available reconstruction algorithms and propose a new one, which is to some extent better adapted to the sampling conditions of such satellites.

1 Introduction

Satellite images are not sampled on an exactly regular grid, but rather on a slightly perturbed grid. The sources of these perturbations include: micro-vibrations of the satellite while it takes the image, and irregularities in the position of the sensors on the image plane. For certain satellite images, the combined effect of these perturbations can be automatically estimated for each image, by different means developed at CNES. Physical models of the satellites exists, which allow to predict certain vibration modes that can be activated at one time or another. For this reason the perturbations in the

*A very early version of this chapter was published in the proceedings of the GRETSI'01 symposium in Toulouse [18].

sampling grid can be modeled quite accurately by:

$$\varepsilon(x) = \sum_{k=1}^M a_k(x) \cos(\langle \omega_k, x \rangle + \varphi_k) \quad (2.1)$$

where the modulation functions $a_k(x)$ are extremely smooth and bounded by a constant A_k which is usually inversely proportional to the frequency ω_k . The vibration frequencies ω_k themselves are usually at least an order of magnitude smaller than the Nyquist cut-off frequency associated to the sampling rate. The number M of vibration modes is also relatively small. This produces very smooth perturbations where the distance between successive samples does not deviate from one pixel by more than about 10%. But accumulated over a large distance the absolute perturbation may reach several pixels. All in all the perturbation is so smooth that it does not produce a visible geometrical distortion of the image, so its correction is not strictly necessary from an aesthetic or presentation point of view. However, such perturbations must be taken into account by algorithms which interpolate these images to obtain sub-pixel accuracy. Such an example is the production of highly accurate digital elevation models from stereo pairs [91], or super-resolution of color images from multi-spectral images using a high resolution panchromatic channel [151]. These applications require image registration with an accuracy even finer than 0.1 pixels in the disparity map. To achieve such an accuracy, the micro-vibrations in the original image sampling must be corrected before registration.

In this chapter we study the problem of resampling the image on a regular grid, given its samples on a perturbed grid and the corresponding perturbation. We note that this perturbation can be obtained with a high level of accuracy from cues given both by gyroscopes mounted on the satellite and by analyzing the images themselves [142]. Nevertheless we shall not deal here with the estimation of the perturbation (*i.e.* the position of the sampling points in the irregular grid), and we shall rather assume that the irregular grid is given with a high level of accuracy.

The chapter is organized as follows. First (section 2) we review the conditions required by irregular sampling theory to make the reconstruction problem well-posed, and we analyze how these conditions apply to satellite imaging systems. Then (section 3) we review some state-of-the-art reconstruction algorithms and point out the characteristics of satellite images that are not exploited by these methods. Next (section 4), we propose a new algorithm, based on a pseudo-inverse iteration, which better exploits these characteristics. Finally, we discuss the results of our simulations (section 5).

2 Existence theorems

2.1 Problem statement.

We assume that the continuous image f before sampling is *band-limited*, i.e. that $f \in L^2(\mathbb{R}^2)$ and $\text{supp}(\hat{f}) \subseteq [-\pi, \pi]^2$. We shall note this space of band-limited functions by PW , for Payley-Wiener. We have seen in chapter 1 that this is not strictly the case in real systems, and that sometimes the reciprocal cell needs to be adapted to the transfer function. Furthermore, even with an adapted reciprocal cell, \hat{f} may have an important part of the energy outside this cell. We shall return later in this section to the question of to what an extent sampling systems satisfy these hypotheses. For the moment we keep it because most of the theory on irregular sampling is developed in this framework, and results are not straightforwardly generalizable to arbitrary spectral domains.

Secondly, we assume that we know the positions of the samples

$$\Lambda = \{\lambda_k\}_{k \in \mathbb{Z}^2} \quad , \quad \lambda_k \in \mathbb{R}^2 \quad (2.2)$$

and the values of f at these points

$$\tilde{s}_k = f(\lambda_k) \quad , \quad k \in \mathbb{Z}^2. \quad (2.3)$$

Then the problem consists of finding the regular samples of f from the vector \tilde{s} of irregular samples

$$s_k = f(k) \quad , \quad k \in \mathbb{Z}^2. \quad (2.4)$$

In the general *irregular sampling* case, the sampling set Λ doesn't need to have any particular structure (beyond a minimal density), and some theorems and algorithms apply to this general case. In satellite imaging systems, however, we shall restrict ourselves to the *perturbed sampling* case, in which the samples

$$\lambda_k = k + \varepsilon(k) \quad , \quad k \in \mathbb{Z}^2 \quad (2.5)$$

are a perturbation of the regular grid, where some additional properties may be assumed about the particular form of the perturbation function ε . Classical results make statements about the *amplitude* of ε , even though a model of the form given in equation (2.1) would be more appropriate in satellite imaging. Here we shall emphasize the fact that ε may be assumed regular (low frequency) with respect to the image f .

Before we present the existence results we should make a precision. We are only interested in *stable reconstructions*. The precise sense of stable reconstructions is slightly

different in different settings, but in general terms, it intends to make sure that the operator mapping $f(\Lambda)$ to $f(\mathbb{Z}^2)$ (or equivalently to f) has a bounded norm, not depending on $f \in PW$.

We divide the remaining of this section in two parts. First we study the perturbed sampling case, where precise bounds on the amplitude of ε are known if we are sampled at the Nyquist rate (section 2.2). Then we discuss the results that are known in the general irregular sampling case (section 2.3). These statements have two differences with respect to the results in section 2.2: First they require over-sampling (in the generalized sense of Beurling-Landau densities); Secondly, when interpreted in terms of perturbed sampling, these results can be rewritten in terms of the *regularity* of the grid perturbation ε , instead of its amplitude.

2.2 Perturbed critical sampling

The main result on the limits of *perturbed sampling* in the one-dimensional case is due to Kadec [104] (see also [175, section 1.10, page 42]) and can be stated as follows:

Theorem 6 (1D Kadec). *If there is a constant c such that*

$$|\varepsilon(k)| = |\lambda_k - k| \leq c < \frac{1}{4} \quad (2.6)$$

then the family $\{e^{i\lambda_k \xi}\}_{k \in \mathbb{Z}}$ forms a Riesz basis of $L^2[-\pi, \pi]$. Hence there exists a stable reconstruction formula of any band-limited function $f \in L^2(\mathbb{R})$ from its irregular samples $f(\lambda_k)$.

Note that in this case stability of the reconstruction is formulated in terms of existence of a Riesz basis in the Fourier domain. The fact that the functions $\varphi_k(\xi) = e^{i\lambda_k \xi}$ form a Riesz basis implies that the operator

$$G : L^2[-\pi, \pi] \rightarrow L^2[-\pi, \pi] \quad G(f) = \sum_k \langle f, \varphi_k \rangle \varphi_k \quad (2.7)$$

is invertible, auto-adjoint and positive. Then the family $\{\tilde{\varphi}_k\}$ defined by $\tilde{\varphi}_k = G^{-1}(\varphi_k)$ is also a Riesz basis biorthogonal to the $\{e^{i\lambda_k \xi}\}_{k \in \mathbb{Z}}$. Finally the reconstruction formula is provided by

$$\hat{f}(\xi) = \sum_k \langle \hat{f}, \varphi_k \rangle \tilde{\varphi}_k(\xi) \quad (2.8)$$

$$= \sum_k f(\lambda_k) \tilde{\varphi}_k(\xi) = \sum_k f(\lambda_k) G^{-1} \varphi_k(\xi) \quad (2.9)$$

$$= G^{-1} \sum_k f(\lambda_k) \varphi_k(\xi), \quad (2.10)$$

where G^{-1} may be computed iteratively by a Neumann series, because $G = a(I - R)$, with $a > 0$ and $\|R\| < 1$. However this iteration may be too slow to converge and more efficient numerical methods are required.

Another important observation about Kadec's theorem is that the maximal value of the constant for the theorem to hold is known exactly. In fact we can build counterexamples with $\varepsilon(x) \leq \frac{1}{4}$, for which the theorem doesn't hold (see [175, p 44 and section 3.3 p 122]). This is not the case for the generalization to two dimensions:

Theorem 7 (2D Kadec). *If the perturbation ε is such that*

$$|\varepsilon(k)| = \|\lambda_k - k\| \leq 0.11 \quad (2.11)$$

then the family $\{e^{i\langle \lambda_k, \xi \rangle}\}_{k \in \mathbb{Z}^2}$ forms a Riesz basis of $L^2([-\pi, \pi]^2)$. Hence there exists a stable reconstruction formula of any band-limited function $f \in L^2(\mathbb{R}^2)$ from its irregular samples $f(\lambda_k)$.

The first generalization of this kind is due to Favier and Zalik [71], where they give a bound of 0.05, and this bound was later improved by Chui and Shi [45]. Nevertheless the bound is not shown to be sharp as in the one-dimensional case.

For a particular kind of perturbations that we call *separable* Jaffard [98] showed that the two-dimensional case can be reduced to the one-dimensional case, and therefore the 0.25 constant is still valid.

Theorem 8 (separable 2D Kadec). *Assume that the perturbation ε is separable, i.e. the perturbation*

$$\varepsilon(n, m) = \begin{pmatrix} \delta_n^1 \\ \delta_n^2 \end{pmatrix} + \begin{pmatrix} 0 \\ \theta_m^2 \end{pmatrix} \quad (2.12)$$

can be expressed as a combination of an arbitrary vector only depending on the horizontal coordinate n , plus a vertical correction only depending on the vertical coordinate m . If in addition there is a constant c such that

$$|\delta_n^1| \leq c \quad \text{and} \quad |\theta_m^2| \leq c \quad \text{and} \quad c < \frac{1}{4} \quad \text{for all } n \in \mathbb{Z} \text{ and } m \in \mathbb{Z} \quad (2.13)$$

then the family $\{e^{i\langle \lambda_k, \xi \rangle}\}_{k \in \mathbb{Z}^2} = \{e^{i[(n+\delta_n^1)\xi_1 + (n+\delta_n^2+m+\theta_m^2)\xi_2]}\}_{(n,m) \in \mathbb{Z}^2}$ forms a Riesz basis of $L^2([-\pi, \pi]^2)$.

Observe in particular that there is no condition on δ_n^2 .

The proof proceeds by showing that the Riesz basis biorthogonal to $\{e^{i\langle \lambda_k, \xi \rangle}\}_{k \in \mathbb{Z}^2}$ can be written as $\{\tilde{\varphi}_n(\xi_1)\tilde{\psi}_m(\xi_2)e^{i\delta_n^2\xi_2}\}_{(n,m) \in \mathbb{Z}^2}$ where $\{\tilde{\varphi}_n(\xi_1)\}_{n \in \mathbb{Z}}$ and $\{\tilde{\psi}_m(\xi_2)\}_{m \in \mathbb{Z}}$ are (one-dimensional) Riesz bases biorthogonal to $\{e^{i(n+\delta_n^1)\xi_1}\}_{n \in \mathbb{Z}}$ and $\{e^{i(m+\theta_m^2)\xi_2}\}_{m \in \mathbb{Z}}$ respectively. Hence the result follows by applying twice the one-dimensional Kadec theorem.

Jaffard's result is of interest in satellite imaging because many of the sampling systems analyzed in chapter 1 actually impose that the perturbations be separable if we choose the coordinate system in a convenient way (see annex B, section 5 for more details). Changing the coordinate system also changes the shape of the canonical reciprocal cell $[-\pi, \pi]^2$. In some cases this change goes in the correct sense, making f more close to band-limited, in other cases the converse is true. In table 2.1, we summarize the results of applying both versions of the two-dimensional Kadec theorem to some of the sampling systems in chapter 1. In some cases the band-limited hypothesis becomes "less true" in other cases the bound on the maximal perturbations that are allowed becomes smaller. Clearly the Hypermode system yields the best results, much better than the Sq2 system which offers roughly the same effective resolution. The reader should be warned to take this comparison with care, since the table only represents the best combination of perturbation bounds and closeness to band-limited that we were able to show with the mathematical results that are available. However, these bounds are by no means necessarily the optimal ones.

2.3 Irregular over-critical sampling

When we consider general irregular sampling sets (not necessarily perturbations of the regular grid), the stable reconstruction condition is usually stated as follows.

Definition 7 (set of (stable) sampling). *A set $\Lambda = \{\lambda_k\}_{k \in \mathbb{Z}^d}$ is called a set of stable sampling or simply set of sampling if there is a constant $C > 0$ such that*

$$\|f\|_{L^2} \leq C \|f(\Lambda)\|_{l^2} \quad (2.14)$$

holds for all band-limited functions $f \in PW$.

A classical result by Beurling and Landau [31, 115], states necessary and sufficient conditions for stable sampling in terms of the *Beurling density*. A simplified version of their results can be written as follows:

$$D(\Lambda) = \liminf_{r \rightarrow \infty} \inf_{x \in \mathbb{R}^d} \frac{\#(B_r(x) \cap \Lambda)}{(r)^d} \quad (2.15)$$

where $B_r(x)$ is a cube of side r and centered at x .

Theorem 9 (Beurling-Landau). *If Λ is a set of stable sampling then $D(\Lambda) \geq 1$. Conversely, if $D(\Lambda) > 1$ then Λ is a set of stable sampling.*

	Deviation from band-limitedness		sup $ \varepsilon $
	$D_{non-sep}^*$	D_{sep}^*	
Hipermode	1.7%	1.7%	0.25p = 0.13c
Sq0	18.0%	18.0%	0.25p = 0.25c
Sq1	6.0%	13.0%	0.11p = 0.08c
Sq2	2.2%	13.2%	0.11p = 0.05c
Hex0	16.3%	16.3%	0.25p = 0.25c
Hex1	2.2%	12.7%	0.11p = 0.06c
Hex2	0.6%	13.5%	0.11p = 0.04c

Table 2.1: Comparison of several systems in terms of “band-limitedness” and maximal perturbations allowed for stable reconstruction. $D_{non-sep}^*$ represents the canonical reciprocal cell $[-\pi, \pi]^2$, in which \hat{f} has to be band-limited to apply the non-separable 2D Kadec theorem 7. D_{sep}^* represents the canonical reciprocal cells $[-\pi, \pi]^2$ in which \hat{f} has to be band-limited to apply the separable 2D Kadec theorem 8. This reciprocal cell is obtained after a suitable change of coordinate system as explained in annex B, section 5. In the separable case we can apply the result of Jaffard with maximal perturbation 0.25 pixels, otherwise we can only use the result of Chui with maximal perturbation 0.11 pixels. (Remember from chapter 1 that this distance in pixels represents different values for the different sampling systems if expressed as a multiple of the size c of the sensors, which represents a constant distance on the ground.) However, for some changes of coordinate system the image becomes very far away from band-limited. Here we measure the deviation from band-limitedness as the percentage $100\sqrt{\int_{\mathbb{R}^2 \setminus D^*} |H|^2 / \int_{D^*} |H|^2}$ of the transfer function H that lies outside of the chosen spectral domain D^* . We observe that except for Hypermode, Sq0 and Hex0, the best band-limitedness occurs in the non-separable case. (However Sq0 and Hex0 are not sufficiently band-limited anyway). So we can apply the separable 2D Kadec theorem, and still obtain the smallest possible deviation from band-limitedness, only for Hypermode, Sq0 and Hex0.

Even if the result was shown originally in the one-dimensional Paley-Wiener framework, several generalizations have been proposed to higher dimensions and shift-invariant L^p spaces [3, 87]. See also [4, p 6] for a review of different generalizations of this result.

Remark 1. The definition of set of sampling should be distinguished from the weaker condition of *set of uniqueness* which simply requires that for all band-limited functions f , if $f(\Lambda) = 0$ then $f = 0$. It can be shown that there are sets of uniqueness with arbitrarily small density whereas sets of sampling satisfy $D(\Lambda) \geq 1$.

Remark 2. Note that Kadec's theorem deals with the critically sampled case where $D(\Lambda) = 1$ which is undetermined by Beurling-Landau's results. On the other hand Beurling's result provides a sufficient condition for stable reconstruction in cases where Kadec's theorem does not apply (because, e.g. the perturbation is too big) but requires oversampling ($D(\Lambda) > 1$) for the sufficient condition to hold.

In practice, even if the Beurling density is larger than 1, the ratio between the frame constants may be extremely large, thus making the problem very ill conditioned. In order to obtain well conditioned systems, practical reconstruction algorithms must impose a stronger condition, that involves a more local version of Beurling's density, namely

Definition 8. A sampling set $X \subseteq \Omega$ is said to be γ -dense if

$$\bigcup_{x \in X} B_\gamma(x) = \Omega. \quad (2.16)$$

Note that γ is, strictly speaking, the inverse of a density, in the sense that in the d -dimensional case

$$D(\Lambda) \geq \frac{1}{\gamma^d}$$

where equality holds e.g. if the sampling set Λ is a regular grid and γ is the minimal radius such that Λ is γ -dense. Hence, to make the link to chapter 1, γ plays the role of spatial resolution R_{nom} , whereas $D(\Lambda)$ plays the role of spectral resolution r_{nom} . In section 3.2 we shall review two algorithms which are shown to be convergent when the sampling set is γ -dense, with $\gamma < 1$ (and thus $D(\Lambda) > 1$). Furthermore Gröchenig [86] proposed an extension of this proof of convergence, which could be used to give a new proof of the Beurling-Landau result, "with the additional benefit of an explicit and efficient numerical algorithm attached to the proof" (quoted from [86]).

3 Available Reconstruction Algorithms

3.1 Based on Kadec's condition (perturbed sampling)

There are a number of works that propose numerical algorithms that are convergent under the conditions of Kadec's theorem [76]. We did not consider such algorithms in this chapter for two reasons: (i) They are formulated in the one-dimensional case, and their generalization to the two-dimensional case is not self-evident, except possibly in the separable case; (ii) They require either high oversampling rates or very small perturbation amplitudes for the filters to be of a reasonable size when the method is generalized to more dimensions. For these reasons they did not seem to be of interest in our particular case.

3.2 Based on Beurling-Landau type conditions (irregular sampling)

ACT algorithm

One of the best performing reconstruction methods available is the ACT algorithm developed initially by Feichtinger, Gröchenig, and Strohmer [72] and further analyzed, refined and generalized by Gröchenig and Strohmer [88], Rauth [141]. The method intelligently combines an accelerated version of the *frame iteration* derived from the proof of Kadec's theorem, with *adaptive weights* in order to improve the condition number of the problem, a *conjugate gradient* iteration which accelerates convergence, and the formulation of the problem as a *Toeplitz system* which makes the computation of each iteration even faster. Furthermore the preparation steps before the conjugate gradient iteration can start, can benefit from the USFFT (for unequally spaced fast Fourier transform) algorithm by Beylkin [32, 33].

More precisely, the algorithm is based on a representation of an $N \times N$ periodic band-limited function f as a trigonometric polynomial of order N^2

$$f(x) = \sum_{n \in [1, N]^2} a_n e^{\frac{2\pi i}{N} \langle n, x \rangle} \quad (2.17)$$

so that the interpolation conditions become

$$\tilde{s}_k = f(\lambda_k) = \sum_{n \in [1, N]^2} a_n e^{\frac{2\pi i}{N} \langle n, \lambda_k \rangle} \quad (2.18)$$

or equivalently in matrix form

$$\tilde{s} = Va, \quad \text{where } V = ((v_{kn})), \quad v_{kn} = e^{\frac{2\pi i}{N} \langle n, \lambda_k \rangle} \quad (2.19)$$

i.e. V is the van der Monde matrix associated to the trigonometric polynomial. Now the problem is reduced to solving for a the system of linear equations (2.19). But if Λ contains some regions with extremely dense sampling the system will not be well balanced. In order to improve the condition number we multiply the k -th equation by a weight

$$w_k = \text{area}(\{x : |x - \lambda_k| < |x - \lambda_j|, \forall j \neq k\}) \quad (2.20)$$

which is inversely proportional to the sampling density at λ_k . In addition, instead of solving the linear system (2.19) directly, it will be more convenient to solve the normal equations as an optimization problem

$$\min_a \|V^* \text{diag}(w)Va - V^* \text{diag}(w)\tilde{s}\|^2 \quad (2.21)$$

because the $N^2 \times N^2$ matrix $T = V^* \text{diag}(w)V$ can be shown to have Toeplitz structure, so that the multiplication Ta can be efficiently computed in $N^2 \log_2(N^2)$ time using Fourier methods. In addition the non-harmonic series

$$t_n = \sum_{k \in [1, M]^2} e^{-\frac{2\pi i}{N} \langle n, \lambda_k \rangle} w_k \quad (2.22)$$

$$b_n = \sum_{k \in [1, M]^2} e^{-\frac{2\pi i}{N} \langle n, \lambda_k \rangle} w_k \tilde{s}_k \quad (2.23)$$

defining $T = V^* \text{diag}(w)V$ and $b = V^* \text{diag}(w)\tilde{s}$ in equation (2.21) can be approximated using the USFFT [32] in $CM^2 \log_2(M^2)$ time each, where C is a constant, which is inversely proportional to the required precision.

Algorithm 1 Reconstruction from irregular samples using ACT algorithm

Require: M^2 irregular samples in vector \tilde{s} , and degree $N^2 \leq M^2$ of trigonometric polynomial.

Ensure: N^2 regular samples in vector s .

- 1: Compute $T = V^* \text{diag}(w)V$ and $b = V^* \text{diag}(w)\tilde{s}$ using the USFFT
 - 2: Minimize $\|Ta - b\|^2$ using conjugate gradients.
 - 3: Compute the regular samples $s_k = f(k)$ for $k \in [1, N]^2$ by applying the inverse FFT to a .
-

The overall procedure can be summarized as shown in algorithm 1. The complexity of step 1 is as we said $O(2M^2 \log(M^2))$, the second step takes $4N^2 \log_2(N^2)$ per iteration. The conjugate gradient algorithm ensures convergence (under certain conditions) in N^2 iterations, but numerically, there is little precision to gain beyond about N iterations. Finally step 3 is computed very fast in $N^2 \log_2(N^2)$ iterations. Note that the

method is particularly efficient for highly oversampled data, *i.e.* $N \ll M$, for two reasons: (i) if we are highly oversampled, as we shall see later the convergence rate of CG is faster; (ii) The bulk of the work depending on the large number M^2 of samples is done in step 1. The rest of the algorithm's complexity depends only on the size N^2 of the spectrum. In our case, however, we are critically sampled ($M = N$) and the total complexity becomes about $4N^3 \log_2(N^2)$, if CG converges in N iterations. In [88, 141] many criteria are given, to obtain estimates of error at each iteration, to decide when to stop the algorithm (because no further improvement of the approximation is possible), as well as to estimate the optimal size N^2 of the spectrum.

The convergence rate of the CG algorithm is determined by the condition number $\kappa = \text{cond}(T)$, or the ratio of the largest to the smallest eigenvalue of T . More precisely at each iteration the approximation error is decreased by a factor $\frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}$ [83]. Gröchenig [85], Gröchenig and Strohmer [88] provided a useful characterization of the condition number of T in the 2-dimensional case

Proposition 3 (ACT convergence rate). *If the sampling set is γ -dense with*

$$\gamma < \frac{\log 2}{4\pi} \quad (2.24)$$

then the condition number of T is

$$\kappa \leq \frac{4}{(2 - e^{4\pi\gamma})^2} \quad (2.25)$$

and algorithm 1 converges to the exact solution s .

Note that the sampling has to be much more dense than the critical sampling rate, for the algorithm to ensure convergence to the exact result, which is not the case in satellite imaging.¹ Nevertheless, even when T is not invertible (and has an infinite condition number), the CG iteration chooses among the minimizers of $\|Ta - b\|$ the one of minimal norm, *i.e.* $a = T^+b$, where T^+ is the pseudo-inverse of T . Therefore, if we know the a-priori spectral decay rate of the image

$$|\hat{f}(\omega)| \leq C\phi(\omega)$$

typically $\phi(\omega) = (1 + \omega)^{-r}$ for $r \in [1, 2]$ or a combination of this natural decay rate, and the known transfer function, in that case we can regularize the solution by imposing this decay rate. Specifically we minimize the modified system

$$\min \|TDe - b\|, \quad a = De \quad (2.26)$$

¹ The bound on the density γ can be significantly relaxed in the one-dimensional case and in the case of a "separable" two-dimensional perturbation.

where $D = \text{diag}(\{\phi(\frac{2\pi}{N}n)\}_n)$ is a diagonal matrix containing the corresponding values of $\phi(\omega)$. This way we shall obtain the solution $a = D(TD)^+b$, which is among all minimizers of $\|Ta - b\|$ the one which best follows the prescribed spectral decay rate, i.e. $D^{-1}a$ is of minimal norm.

This discussion explains why the ACT algorithm of Gröchenig and Strohmer [88] provides good approximations to the exact solution, even when the convergence conditions in proposition 3 are not satisfied as in the case of satellite imaging, while at the same time the convergence rate is not that good.

We also tried limiting the size N of the spectrum to a smaller value than M , in such a way that the density condition in proposition 3 is satisfied. In this case we cannot expect to recover the high frequencies of f , but convergence should be faster. In practice, the accelerated convergence is minimal and does not compensate for the loss in accuracy, so we always kept the regularized version described above, imposing a polynomial decay for \hat{f} .

Aldroubi's algorithm

More recently (actually we discovered this method by the end of this project) Aldroubi and Gröchenig [4] proposed a different method which has different convergence properties, and is directly formulated in the more general shift-invariant weighted d -dimensional L^p spaces $V_\nu^p(\phi)$, where ν is a weight that is introduced in the L^p norm to control the decay rate of the function in the spatial domain, and the generator ϕ must satisfy certain regularity properties that depend on p and ν . The particular Payley-Wiener framework we are dealing with here is obtained for $p = 2$, $\nu = 1$ and $\phi = \text{sinc}$.

Assume that our sampling set Λ is γ_0 -dense, so that $\mathbb{R}^d = \bigcup_k B_\gamma(\lambda_k)$ for any $\gamma > \gamma_0$. Now build a partition of unity $\{\beta_k\}_k$ such that

$$\begin{aligned} 0 &\leq \beta_k \leq 1 \\ \text{supp}(\beta_k) &\subset B_\gamma(\lambda_k) \\ \sum_k \beta_k &= 1. \end{aligned}$$

Such a partition is called a *bounded partition of unity*. Then define an operator Q_Λ by :

$$Q_\Lambda f = \sum_k f(\lambda_k)\beta_k. \quad (2.27)$$

Note that for the particular partition of unity given by $\beta_k = \mathbb{1}_{V_k}$ where V_k is the Voronoi cell associated to λ_k in the set Λ , the quasi-interpolant operator Q_Λ corresponds to nearest neighbour interpolation. This construction, however, allows for a more smooth and

bounded interpolation in the case where we take for instance irregular spline functions of higher order.

Furthermore we consider the projection operator P from L_v^p to $V_v^p(\phi)$ (in the Payley-Wiener framework this corresponds to removing the frequencies beyond $[-\pi, \pi]^2$), then for sufficiently small γ it is shown that the operator $PQ_\Lambda - I$ is a contraction. This justifies the iteration

$$\begin{cases} f_1 &= PQ_\Lambda f \\ f_{n+1} &= PQ_\Lambda(f - f_n) + f_n = f_1 + (PQ_\Lambda - I)f_n \end{cases} \quad (2.28)$$

with convergence

$$\|f - f_n\|_{L_v^p} \leq C\alpha^n \|f - f_1\|_{L_v^p} \quad (2.29)$$

for some convergence rate $\alpha < 1$ depending on γ .

Now we explain how we implemented the operators PQ_Λ and $(PQ_\Lambda - I)$ in the Payley-Wiener framework, since this is not explained in detail in [4]. At each step we assume that f_n is represented by its samples on \mathbb{Z}^2 . This is sufficient to completely reconstruct $f_n(x)$ at any point $x \in \mathbb{R}^2$, because –thanks to the projector P – we have $f_n \in PW$. On the other hand, the result of applying the quasi interpolator $Q_\Lambda f_n$ or $Q_\Lambda f$ must be computed on a finer grid $(\mathbb{Z}/z)^2$ for some zoom factor $z > 1$, otherwise there would be no sense in applying the projector P afterwards. The choice of this zoom factor is for the moment quite arbitrary, in the absence of a criterion that allows us to compute its influence in the approximation error.² As a rule we took $z = n$ when a precision of 10^{-n} was required. This choice seemed to be enough in our experiments (larger values of z did not significantly improve the convergence rate), but it is possibly excessive.

In order to save time in the calculation of the nearest-neighbour interpolation, we pre-compute at the beginning of the iteration the mapping $x \mapsto k(x)$ which assigns to each point $x \in (\mathbb{Z}/z)^2$ in the fine grid, the index k to the closest point $\lambda_{k(x)} \in \Lambda$ in the irregular sampling grid. Then the quasi-interpolator is computed for all $x \in (\mathbb{Z}/z)^2$ simply by

$$(Q_\Lambda f)(x) = \tilde{s}_{k(x)} \quad (2.30)$$

$$(Q_\Lambda f_n)(x) = f_n(\lambda_{k(x)}) \quad (2.31)$$

² $Q_\Lambda f$ has no reason to be band-limited, nor in any spline-like space, even if f is band-limited. However the Fourier coefficients of $Q_\Lambda f$ should exhibit a certain decay rate of its Fourier coefficients, which depends on the sampling density. From this decay rate it should be possible to compute a more accurate criterion to choose the appropriate zoom factor z for a given precision.

which shows the need of resampling $f_n(\Lambda)$ on the irregular grid from its regular samples $f(\mathbb{Z}^2)$.

This is an easier problem than the inverse irregular to regular sampling problem we are trying to solve, since we only need to do Shannon interpolation. However, we cannot do it simply by Fourier methods, because the target grid is irregular. Separable filtering doesn't work either, unless the irregular grid is separable as well. Therefore, to efficiently solve the problem in the general case, with any given desired precision, we used high-order B-spline interpolation with oversampling, as described by Unser [171]. Briefly, we approximate the sinc filter by the cardinal B-spline:

$$\text{sinc} \approx \beta_{\text{int}}^{(k)} = \mathcal{F}^{-1} \left(\frac{\hat{\beta}^{(k)}}{\sum_{l=-\infty}^{+\infty} \hat{\beta}^{(k)}(w + 2\pi l)} \right) \quad (2.32)$$

where the k -th order B-spline is defined by the recursion $\beta^{(k)} = \beta^{(0)} * \beta^{(k-1)}$, $\beta^{(0)} = \mathbb{I}_{[-\frac{1}{2}, \frac{1}{2}]}$. Note that this has the advantage that the filter can be decomposed in two parts. The denominator is a prefilter with infinite spatial support that is applied in the Fourier domain, where the image has finite support. Then the numerator, which has a small spatial support (of $k \times k$ pixels in the two-dimensional case), is applied in the spatial domain in a non-separable way, since it has to be evaluated on the irregular grid.

Furthermore, in order to improve accuracy with a minimal computational cost, we use a B-spline approximation of the sinc filter corresponding to o times the original sampling rate (we can do so because f and f_n are band-limited), and we apply it to an o -times oversampled version of the image (o times zero-padded in the Fourier domain). This improves accuracy because the B-spline approximation is very good except close to the Nyquist frequency. With oversampling, the part of the spectrum where the B-spline approximation is bad, lies outside the spectral support of the image.

From the spline interpolation error kernel in Unser's article, we can compute the expected error or the worst-case error given the spectral behaviour of f . Figure 2.1 shows the expected relative L^2 error when interpolating a band-limited function f by k -th order B-spline interpolation with o -times oversampling, for different orders and oversampling factors. The graphic shows also the number of flops per pixel needed to interpolate for a given order and oversampling factor. Given this pre-computed chart, and a desired approximation precision δ , we simply choose from it the value of k and o which minimizes the computational cost and with expected error smaller than δ .

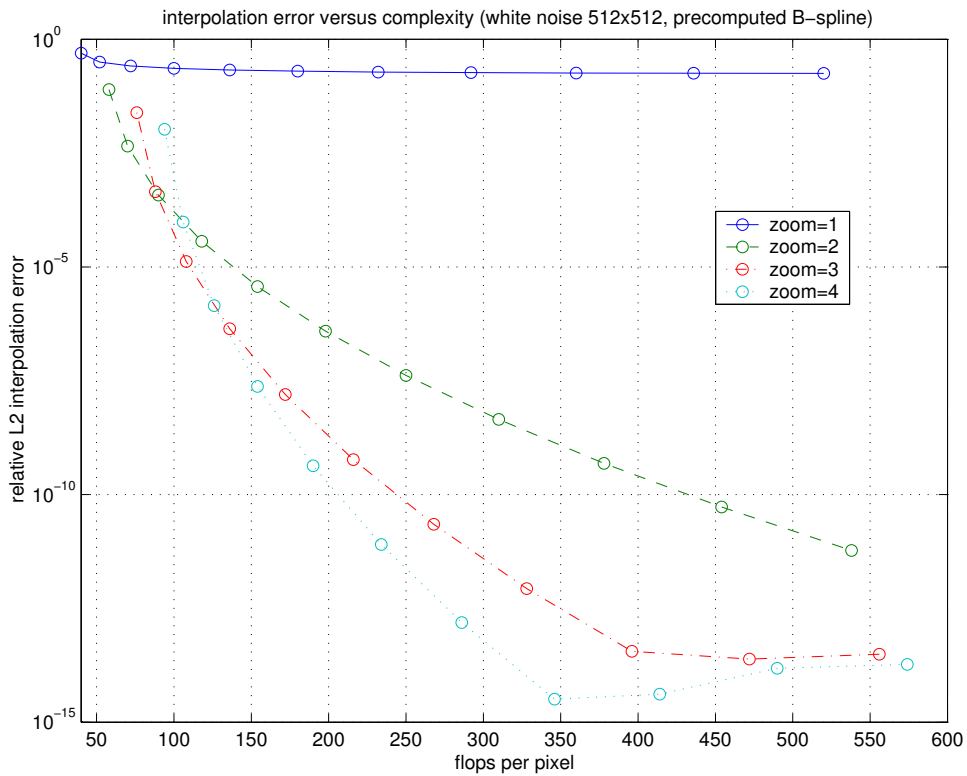


Figure 2.1: *Expected error in spline interpolation versus computational complexity.* The errors are computed on the basis of a white-noise image, with spectrum limited to $[-\pi, \pi]^2$. The computational complexity is computed on the base of an image of size 512×512 , and assuming that the values of the B-spline functions are tabulated to the required accuracy. Both the error and number of operations per interpolated pixel are computed for o -oversampled k -th order B-spline interpolation with orders $k \in \{1, 3, 5, \dots, 21\}$, and oversampling factors $o \in \{1, 2, 3, 4\}$. Beyond precision 10^{-14} results are inaccurate because we are close to the machine precision.

4 A new reconstruction algorithm

The slow convergence of Gröchenig's algorithm for our particular kind of data and the absence of a convergence theorem in this situation led us to the study of new algorithms. In particular we tried to combine : (i) the fact that we are in the *perturbed* sampling case (which is only exploited by the algorithms in section 3.1), with (ii) the fact that the perturbation is very *smooth* (which is only exploited by the algorithms in section 3.2), and (iii) the fact that the image itself is very *regular*, since it may have a fast decay near the Nyquist frequency, due to the instrument's transfer function and higher sampling rates than conventional systems.

The algorithm is very similar to Aldroubi's algorithm, it is based on alternating two operators, whose combination is a perturbation of the identity, and such that one of them simulates the original perturbed sampling. Only the choice of operators is different.

Observe that the irregular sampling process may be written in the following manner. Given a band-limited image f (or equivalently its regular samples $s = \Delta_{\mathbb{Z}^2} \cdot f$, since we can reconstruct $f = \text{sinc} * s$), and the perturbation ε , consider the warping function $\phi_\varepsilon = I + \varepsilon$. Then the irregular samples can be written as a regular sampling of the warped image

$$\tilde{s} = \Delta_{\mathbb{Z}^2} \cdot (f \circ \phi_\varepsilon) = \Delta_{\mathbb{Z}^2} \cdot ((\text{sinc} * s) \circ \phi_\varepsilon) = A^+ s. \quad (2.33)$$

We abbreviate this irregular sampling operator by A^+ .³ Now all the problem is reduced to inverting the operator A^+ . Its inverse is still a linear operator, but not necessarily a convolution-sampling operator. Nevertheless we can still approximate the inverse of A^+ by a convolution-sampling pair (that we shall call A^-) as follows:⁴

$$s_1 = \Delta_{\mathbb{Z}^2} \cdot ((\text{sinc} * \tilde{s}) \circ \phi_\varepsilon^{-1}) = A^- s. \quad (2.34)$$

This corresponds to performing the inverse warp from the irregular samples, seen as regular samples of the warped image $\tilde{f} = f \circ \phi_\varepsilon$, by assuming that \tilde{f} is also band-limited to $[-\pi, \pi]^2$, which is not the case. In addition, some information has been lost because \tilde{f} has been sampled at the Nyquist rate corresponding to $[-\pi, \pi]^2$, while the spectrum of \tilde{f} is usually larger.

³ Note that this operator can be efficiently calculated using the FFT, spatial convolutions and sampling, by the same oversampled B-spline procedure described at the end of the previous section.

⁴Note that the computation of $\phi_\varepsilon^{-1}(\mathbb{Z}^2)$ can also be formulated as a perturbed sampling problem, but in this case, we largely have the necessary oversampling rate. Hence we can compute it very quickly using the ACT algorithm.

Note that in the case of constant ε , both operators A^+ and A^- reduce to (opposite) translations by Shannon interpolation and one is the inverse of the other. Therefore, for very smooth “almost constant” ε we could expect A^- to be close to the inverse of A^+ .

In the following we shall assume that A^- is a good approximation of the inverse of A^+ in the sense that

$$A^-A^+ = (I - \alpha), \quad \text{with } \|\alpha f\| \leq a\|f\|, \text{ and } a < 1 \quad (2.35)$$

for all f in a certain class. Then we have the equality

$$s_1 = A^-A^+s = (I - \alpha)s \quad (2.36)$$

And $(I - \alpha)^{-1}$ can be expanded as a Neumann series, to recover s from s_1 by the following iteration

$$s_n = s_1 + \alpha s_{n-1} = (I - \alpha^n)s \quad (2.37)$$

so that

$$\|s_n - s\| = \|\alpha^n s\| \leq a^n \|s\| \quad (2.38)$$

Under the assumption given in equation (2.35), this means that the approximation error is reduced at each iteration by a factor $a < 1$, so the convergence is geometric.

In practice the operator α is applied in two steps: First, a simulation $\tilde{s}^{(n-1)}$ of the perturbed image from the current estimate $s^{(n-1)}$ of the regular image:

$$\tilde{s}^{(n-1)} = A^+s^{(n-1)}; \quad (2.39)$$

and secondly a correction of the errors found in this simulation with respect to the known perturbed image \tilde{s} :

$$s^{(n)} = s^{(n-1)} - A^-(\tilde{s}^{(n-1)} - \tilde{s}). \quad (2.40)$$

4.1 Convergence analysis

Now we return to the conjecture in equation (2.35). It turns out that we cannot assume $\|\alpha\| < 1$ in general, but we may obtain $\|\alpha f\| < \|f\|$ for all sufficiently regular f , if ε is sufficiently small in amplitude:

Proposition 4. *Let $f \in L^2(\mathbb{R}^2)$ with $\text{supp}(\hat{f}) \subseteq [-\pi, \pi]^2 = R$, and let $s = \Delta_{\mathbb{Z}^2} \cdot f$ be the regular samples of f , and $\tilde{s} = \Delta_{\mathbb{Z}^2} \cdot (f \circ \phi_\varepsilon)$ its irregular samples, with a small perturbation ε*

with a much smaller spectral support $\text{supp}(\hat{\varepsilon}) \subset R$. Then the operator α of the pseudo-inverse algorithm can be expanded up to third order as follows:

$$\alpha s = \Delta_{\mathbb{Z}^2} \cdot \left(\langle [\mathcal{F}^{-1}(\mathbb{1}_{R^c}) * D \langle Df, \varepsilon \rangle - D \text{alias}_R \langle Df, \varepsilon \rangle], \varepsilon \rangle \right) + O(\varepsilon^3) \quad (2.41)$$

The proof is given in appendix A, at the end of this chapter. Concerning the notation used in this proposition, observe that the aliasing operator associated to a sampling grid Γ (in our case $\Gamma = \mathbb{Z}^2$) or equivalently, to its Voronoi reciprocal cell R is

$$(\text{alias}_R f)(x) = \mathcal{F}^{-1}(\mathbb{1}_R) * \left(\sum_{k \in \Gamma^*, k \neq 0} f(x) e^{i \langle k, x \rangle} \right) \quad (2.42)$$

or equivalently in the Fourier domain

$$\mathcal{F}(\text{alias}_R f)(\omega) = \mathbb{1}_R \cdot \left(\sum_{k \in \Gamma^*, k \neq 0} \hat{f}(\omega + k) \right) \quad (2.43)$$

i.e. the error that is incurred when sampling and sinc-interpolating a non-band-limited function f .

Observe that αs doesn't have a linear term, it depends quadratically on ε . From this expression we can also see that $\|\alpha s\|$ depends essentially on the part of the spectrum of $\langle Df, \varepsilon \rangle$ which goes beyond our spectral domain R . But this can be controlled if both ε has a small spectral support, and f has relatively weak high frequencies.

Proposition 5. *If ε is of the form given in equation (2.1) then*

$$\|\alpha s\|_2 \leq 4\pi^2 \|\varepsilon\|_{1,\infty} \left(\sum_{k=1}^M \nu_k \|A_k\|_1 (\|\hat{f} \mathbb{1}_{B_k^+}\|_2 + \|\hat{f} \mathbb{1}_{B_k^-}\|_2) \right) + O(\varepsilon^3), \quad (2.44)$$

where $\|A_k\|_1$ denotes the sum of the amplitudes in the x and y directions, $\|\varepsilon\|_{1,\infty} = \sup_{x \in \mathbb{R}^2} (|\varepsilon_1(x)| + |\varepsilon_2(x)|)$ is the maximal 1-norm of ε in the whole image, and $\nu_k = (1 + \|\omega_k\|_\infty / (2\pi))$ only varies between 1 (for low frequencies) and 1.5 (for high frequencies). Finally the spectral domain B_k^+ is

$$B_k^+ = \{\omega \in R : \exists \gamma \in 2\pi\mathbb{Z}^2, \gamma \neq 0, (\omega + \omega_k + \gamma) \in R\} \quad (2.45)$$

the high frequency ring of width ω_k within $R = [-\pi, \pi]^2$. Similarly B_k^- is a symmetrical set with respect to the origin.

If in addition we assume that f is sufficiently regular, in the sense that

$$\|\hat{f}(w) \mathbb{1}_{B_k^+}(w)\| \leq \frac{\text{area}(B_k^+)}{\text{area}([-\pi, \pi]^2)} \|\hat{f}\| \quad (2.46)$$

then the upper bound (2.44) can be simplified to

$$\|\alpha s\|_2 \leq 4\pi^2 \|\varepsilon\|_{1,\infty} \left(\sum_{k=1}^M \mu_k \|A_k\|_1 \right) + O(\varepsilon^3), \quad (2.47)$$

where $\mu_k = \left(1 + \frac{\|\omega_k\|_\infty}{2\pi}\right) \frac{2\|\omega_k\|_\infty}{\pi}$ varies quadratically between 0 and 3, depending on the frequency ω_k .

The proof is given in appendix A, at the end of this chapter.

This proposition shows clearly that α may in general be unbounded. The only way to ensure that $\|\alpha f\| \leq a\|f\|$ with $a < 1$ is to ensure that f has relatively small high frequencies. This is usually the case for the initial image f , because of the acquisition conditions. But it may not be the case for the iterates f_n . Thus, a way to avoid the divergence of this algorithm is to precede each application of α by a projector P which controls the high frequencies, for instance:

$$Pf = \mathbb{1}_{R \setminus B_{k_0}} * f, \quad \text{where } k_0 = \arg \max_k \|\omega_k\|_\infty, \quad B_k = B_k^+ \cup B_k^- \quad (2.48)$$

To complete the analysis we consider the complementary projector

$$Qf = \mathbb{1}_{B_{k_0}} * f \quad (2.49)$$

which satisfies $P + Q = I$ within the Payley-Wiener space. Now it is more reasonable to assume that $\|\alpha P\| \leq a < 1$. The modified iteration is:

$$s_1 = A^- A^+ s = (I - \alpha P)s - \alpha Qs \quad (2.50)$$

$$s_n = s_1 + \alpha P s_{n-1} = (I - (\alpha P)^n)s - \left(\sum_{k=0}^{n-1} (\alpha P)^k \right) \alpha Qs \quad (2.51)$$

so that

$$\frac{\|s_n - s\|}{\|s\|} \leq a^n + \left(\frac{1 + a^n}{1 - a} \right) \frac{\|\alpha Qs\|}{\|s\|} \quad (2.52)$$

This equation shows that the modified iteration has a limit to accuracy which is determined by $\|\alpha Qs\|$. This term will be small when the high frequencies Qs of s are relatively small, and α does not excessively amplify these high frequencies, which is the case when ε is smooth too.

Now we return to the problem of showing that $\|\alpha P\| \leq a < 1$. The main idea is to try to deduce this from proposition 5. Now a major problem in applying it is that ε may have to be small, so that the unknown third order term $O(\varepsilon^3)$ is smaller than a given constant $c < 1$. To solve this problem we have to apply this result locally, in such

a way that the amplitude of ε is sufficiently reduced.

More precisely, assume that within any interval I of size smaller than L we have that $\max_{x,y \in I} |\varepsilon(x) - \varepsilon(y)| \leq 2\delta$. Then we can consider $\tilde{\varepsilon} = \varepsilon - \bar{\varepsilon}$, which satisfies $|\tilde{\varepsilon}| \leq \delta$ and hence in proposition 5, the third order term is smaller than $c < 1$. This analysis is valid because adding a constant to ε doesn't change α , which can be easily verified.

Now we still have to translate a statement about a global upper bound of $\|\alpha P\|$, into statements about upper bounds of $\|\alpha P \mathbb{1}_I\|$ for intervals of size smaller than L . We do so by considering a partition of unity:

$$\sum_{k \in \mathbb{Z}^2} g_k = 1, \quad \text{with } \text{supp}(g_k) \subseteq B_L(Lk) \quad (2.53)$$

such as *e.g.* the rescaled first order B-splines $g_k(x) = \beta^{(1)}(x/L - k)$. Then, since α is linear we can write

$$\|\alpha Pf\|^2 = \left\| \sum_{k \in \mathbb{Z}^2} \alpha g_k Pf \right\|^2 = \sum_{m,n \in \mathbb{Z}^2} \langle \alpha g_m Pf, \alpha g_n Pf \rangle \quad (2.54)$$

Suppose that we can bound this term by

$$\sum_{m,n \in \mathbb{Z}^2} \langle \alpha g_m Pf, \alpha g_n Pf \rangle \leq \sum_{m,n \in \mathbb{Z}^2} b_{n,m} \|g_m Pf\| \|g_n Pf\| \quad (2.55)$$

where for all n , we have $\sum_{m=-\infty}^{\infty} |b_{n,m}| \leq a^2 < 1$.⁵ If this conjecture holds, then we have

$$\|\alpha Pf\|^2 \leq y^T B y \quad (2.56)$$

where $B = ((|b_{n,m}|))$ is the matrix composed of the coefficients $b_{n,m}$ in equation (2.55), and the vector y has coordinates $y_n = \|g_n Pf\|$. Using Gerschgorin's theorem [55], we deduce that the eigenvalues of B lie within the union of the circles D_n , plus the limits $D_{\pm\infty}$, where each D_n is centered at $|b_{n,n}|$ and has radius $r_n = \sum_{m \neq n} |b_{m,n}|$. Hence all the eigenvalues of B are smaller than $\sum_{m=-\infty}^{\infty} |b_{n,m}| \leq a^2 < 1$. In conclusion $y^T B y \leq a^2 \|y\|^2$, *i.e.*

$$\|\alpha Pf\|^2 \leq a^2 \sum_n \|g_n Pf\|^2 = a^2 \int_{\mathbb{R}^2} |Pf|^2 \left(\sum_n g_n^2 \right) \quad (2.57)$$

⁵ Observe that for $n = m$ the value of $b_{n,m}$ is provided by proposition 5 and the previous discussion on the local application of this proposition. In the case $|n - m| > 1$, the two window functions g_n and g_m have no common support, and the only reason for $\langle \alpha g_m Pf, \alpha g_n Pf \rangle$ to be non-zero is that αf expands the spatial support of f , due to the sinc filters it contains. Finally, the case $|n - m| = 1$ may require a stationarity hypothesis on f in order to show that for a given n , $\sum_{|n-m|=1} |b_{n,m}| \leq 8|b_{n,n}|$.

But since g_n is a partition of unity with $g_n \in [0, 1]$, we have that $g_n^2 \leq g_n$ and hence $\sum_n g_n^2 \leq 1$. In fact this sum is a function which oscillates between 1 and 0.5. Finally, the projection operator can only make decrease the norm of f , and we have:

$$\|\alpha P f\|^2 \leq a^2 \int_{\mathbb{R}^2} |P f|^2 = a^2 \|P f\|^2 \leq a^2 \|f\|^2. \quad (2.58)$$

In order to complete this proof we still need to show conjecture (2.55) and provide a criterion to compute the maximal size L of a neighborhood, within which ε is sufficiently small, for the third order term $O(\varepsilon^3)$ in proposition 5 to be smaller than a given constant $c < 1$.

4.2 Numerical approximation

When applying the operator α in equation (2.37) or the operator αP in the modified algorithm in equation (2.51) the operators A^+ and A^- will be approximated in our algorithm by $\alpha_\delta = A_\delta^- A_\delta^+ - I$, whereas in equation (2.36) it will be approximated by $\alpha'_\delta = A_\delta^- A^+ - I$, since the first application of A^+ is performed by the acquisition system, not by our algorithm. It is easy to show that $\lim_{\delta \rightarrow 0} \alpha_\delta = \lim_{\delta \rightarrow 0} \alpha'_\delta = \alpha$ where $\delta = 1/\sigma$ for the damped sinc approximation, and $\delta = 1/m$ for the cardinal B-spline approximation. Hence for a sufficiently small δ we shall have $\|\alpha_\delta\| < a < 1$ and $\|\alpha'_\delta\| < a < 1$. To summarize, the iterations with the numerical approximation of the sinc filter are as follows:

$$s_\delta^{(1)} = (I - \alpha'_\delta) s \quad (2.59)$$

$$\begin{aligned} s_\delta^{(n)} &= \alpha_\delta \cdot s_\delta^{(n-1)} + s_\delta^{(1)} \\ &= \left(I + \left(\sum_{i=1}^{n-2} \alpha_\delta^i \right) (\alpha_\delta - \alpha'_\delta) - \alpha_\delta^{n-1} \alpha'_\delta \right) \cdot s \end{aligned} \quad (2.60)$$

Observe that the relative error $\|s_\delta^{(n)} - s\|/\|s\|$ still has a term $\|\alpha_\delta^{n-1} \alpha'_\delta\| < a^n$ that decreases geometrically with the iterations, but a second term appears which may increase with the iterations but is bounded by $\frac{a}{1-a} \|\alpha_\delta - \alpha'_\delta\|$ (provided $\|\alpha\| \leq a < 1$). This second term can be kept below the noise level as long as $a < 1$ and δ is sufficiently small.

5 Experiments

In order to test and compare the performance of the different algorithms we constructed a set of simulated satellite images. We started from a very high resolution

aerial image g which was later filtered with the transfer function of different imaging systems, particularly a CCD array with sharp optics, and the supermode and hypermode systems (see [116, 148] and annex B). In all cases the transfer function was truncated beyond the corresponding Voronoi cell so that aliasing effects were not taken into account, and we obtained perfectly band-limited images. The application of the transfer function is done simply to simulate the spectral contents of images obtained with different instruments. This produced the reference image s with spectral support $[-\pi, \pi]^2$.

Then we simulated different perturbations ε such that $\text{supp}(\hat{\varepsilon}) \subseteq [-\frac{\pi}{T}, \frac{\pi}{T}]^2$ for different values of $T > 1$, and for different standard deviations for $\varepsilon(x)$ (typically 0.25, 0.5 and 1). With the given perturbation we simulated the perturbed images \tilde{s} with a high precision (usually 10^{-8}) using the oversampled B-spline interpolation technique described at the end of section 3.2 (Aldroubi's algorithm). Finally we added some white noise to the irregular samples with standard deviation 10^{-3} times smaller than the standard deviation of the image ($SNR = 60dB = 10^3$). Then we asked the three methods to recover s up to a precision of 10^{-3} .

Concerning the pseudo-inverse method, we observed that the stabilized version (the one in equations (2.50) and (2.51) which applies αP at each iteration), even if it avoids divergence, it converges with an accuracy which is less precise than the one obtained by intermediate iterates of the divergent version in equations (2.36) and (2.37). Furthermore, the point at which it starts to diverge can be detected (without knowledge of s of course), at the point where the residual stops decreasing. This usually happens one or two iterations after the actual error stops decreasing. For this reason we always applied the non-convergent version, because it usually provides a better accuracy, and the divergence point can be determined. In the graphs, however, we show many more iterations beyond the detection of the divergence point (for all algorithms), in order to show its behaviour.


Figures 2.2 to 2.5 show some representative results. The experiments confirm the influence of the regularity of the image and of the the amplitude and the regularity of the perturbation in the convergence of the pseudo-inverse algorithm. They also show how the ACT algorithm, usually provides the finest accuracy, but at the cost of more computations. The pseudo-inverse algorithm, on the other hand, is the fastest to provide a good approximation, especially for regular images and perturbations, even though its precision may be less fine than that of the other two methods if the image is very sharp. In many cases the computational effort of the pseudo-inverse with re-

spect to the ACT algorithm to reach the same accuracy is a factor of 3, 5 or even 10 times smaller. Finally, Aldroubi's method presents an intermediate behaviour, often producing an approximation level just between those of the pseudo-inverse and the ACT algorithms, and also with an intermediate computational cost.

Further experiments on the pseudo-inverse algorithm confirm the convergence properties described in section 4.1. For instance, if the image has a spectral contents similar to Hipermode, the perturbation consists of a single sinusoidal oscillation with amplitude A and period T , then the minimal reconstruction error $\|f - f_n\|$ during the pseudo-inverse iterations was at least 5 times smaller than the initial error $\|f - \tilde{f}\|$ whenever $\frac{A}{T} < 0.1$. Similarly, if the perturbation was a colored noise with standard deviation A (maximal perturbation is more than $2.5A$) and spectrum in $[-\frac{2\pi}{T}, \frac{2\pi}{T}]^2$, then we obtained a reconstruction error 5 times smaller than the initial error whenever $\frac{A}{T} < 0.05$. This means that if the image is sufficiently smooth, for perturbations with frequencies smaller than $\frac{\pi}{10}$ for instance we can still obtain a reasonable reconstruction with perturbation amplitudes up to about 2 pixels.


T	5		10		20	
$A \setminus \varepsilon$	sin	rand	sin	rand	sin	rand
0.125	4.1e-09, 6.0e-08	1.4e-04, 3.7e-03	4.3e-09, 6.1e-08	1.7e-04, 4.4e-03	4.3e-09, 6.2e-08	1.5e-06, 3.6e-05
0.250	4.5e-07, 1.3e-05	1.1e-01, 5.7e+00	2.0e-07, 5.7e-06	1.5e-02, 7.3e-01	1.4e-07, 4.2e-06	1.9e-03, 9.4e-02
0.500	1.7e-01, 9.9e+00	1.2e+00, 1.2e+02	3.7e-03, 2.1e-01	7.6e-02, 7.5e+00	2.1e-03, 1.2e-01	2.3e-02, 2.2e+00
1.000	1.9e+00, 2.1e+02	3.2e+00, 5.8e+02	1.7e-01, 1.9e+01	9.0e-01, 1.6e+02	6.9e-02, 7.6e+00	4.0e-02, 7.0e+00
2.000	3.1e+00, 6.2e+02	2.7e+00, 7.7e+02	1.8e+00, 3.5e+02	2.5e+00, 6.9e+02	9.3e-02, 1.9e+01	5.9e-01, 1.7e+02
4.000	2.0e+00, 6.2e+02	2.0e+00, 7.8e+02	2.1e+00, 6.6e+02	2.0e+00, 7.6e+02	1.1e+00, 3.2e+02	2.0e+00, 7.4e+02

Table 2.2: Accuracy test of the pseudo-inverse algorithm. In this experiment the reference image f has spectral contents similar to Hipermode. The perturbation ε is either a purely horizontal oscillation $A \sin(2\pi x/T)$ (sin column) or its Fourier transform is a white noise with standard deviation A inside $[-\frac{2\pi}{T}, \frac{2\pi}{T}]^2$ and zero outside this cell (rand column). Then starting from the perturbed image \tilde{f} , for a given value of T and A , we computed the pseudo-inverse reconstruction to obtain a sequence f_n of approximations to the reference f , and kept the iteration n that minimizes the reconstruction error $\|f_n - f\|$. To the left of each column we show $\frac{\|f_n - f\|}{\|\tilde{f} - f\|}$ (the improvement with respect to the original), and to the right of each column we show $\frac{\|f_n - f\|}{\delta}$ where $\delta = 10^{-3}$ is the approximation level used to chose appropriate B-spline filters.



(a) Image spectrum similar to a CCD array

(b) Image spectrum similar to Supermode



(c) Image spectrum similar to Hypermode

Figure 2.2: The influence of the spectral contents of the image on the performance of the methods. Observe that for images with weak high frequencies such as hypermode, the pseudo-inverse method reaches the attainable precision much faster than the other two methods. For sharper images, however, the pseudo-inverse attains a less precise accuracy than the other two methods. In all cases the method of Gröchenig, Strohmer and Rauth attains the most precise accuracy faster, but intermediate precisions may be attained faster by Aldroubi's method.

In all cases, the perturbation is a colored noise with standard deviation of 0.5 pixels (maximal perturbation is 1.25 pixels), and spectral contents inside $[-\frac{\pi}{T}, \frac{\pi}{T}]^2$ for $T = 20$, and the SNR of the sampled image is 60 dB, *i.e.* the L^2 norm of the noise is 10^{-3} times smaller than the L^2 norm of the image. The approximation level of the oversampled B-spline filters was fixed at the same level.


(a) $\text{supp}(\hat{\varepsilon}) \subseteq (-\frac{\pi}{20}, \frac{\pi}{20})^2$ (b) $\text{supp}(\hat{\varepsilon}) \subseteq (-\frac{\pi}{10}, \frac{\pi}{10})^2$

Figure 2.3: *The influence of the spectral contents of the perturbation on the performance of the methods. Observe that when the frequency of the perturbation increases the performance of the proposed pseudo-inverse method is affected to a larger extent than the other methods. In both cases, the spectral contents of the image is similar to a CCD with good optics. For less sharp images like supermode or hypermode the degradation of the performance of the pseudo-inverse is less important.*

In both cases the SNR of the sampled image is 60 dB, *i.e.* the L^2 norm of the noise is 10^{-3} times smaller than the L^2 norm of the image. The approximation level of the oversampled B-spline filters was fixed at the same level. The perturbation is a colored noise with standard deviation of 0.5 pixels (maximal perturbation is 1.25 pixels), and spectral contents inside $[-\frac{\pi}{T}, \frac{\pi}{T}]^2$ for $T = 20$ (subfigure a), or $T = 10$ (subfigure b).

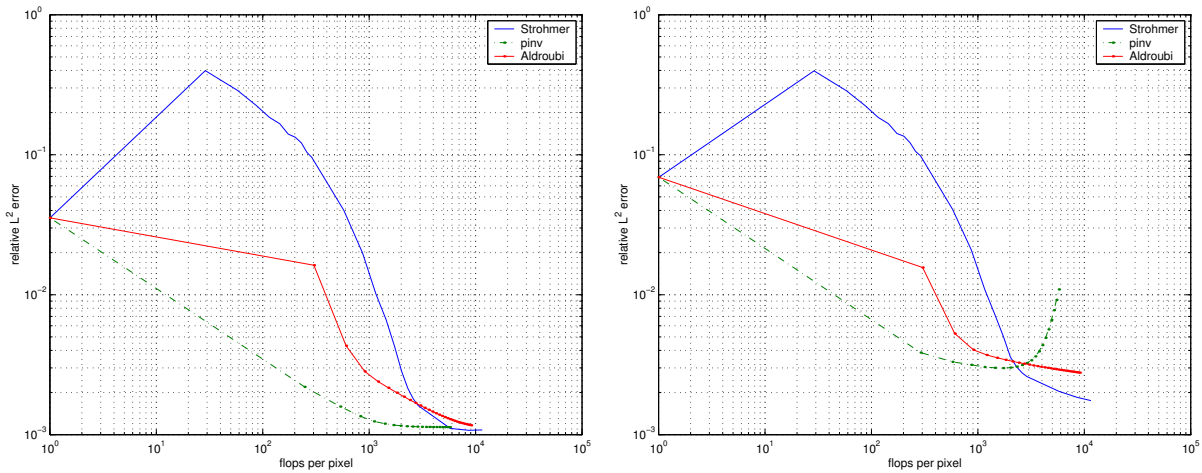

(a) $\text{Var}\{\varepsilon\} = 0.25^2$ (b) $\text{Var}\{\varepsilon\} = 0.5^2$ (c) $\text{Var}\{\varepsilon\} = 1$

Figure 2.4: The influence of the amplitude of the perturbation on the performance of the methods. Observe that the attainable accuracy is affected for all methods by this amplitude.

In all cases, the spectral contents of the image is similar to supermode and the SNR of the sampled image is 60 dB, *i.e.* the L^2 norm of the noise is 10^{-3} times smaller than the L^2 norm of the image. The approximation level of the oversampled B-spline filters was fixed at the same level. The perturbation is a colored noise with standard deviation of 0.5 pixels (maximal perturbation is 1.25 pixels), and spectral contents inside $[-\frac{\pi}{T}, \frac{\pi}{T}]^2$ for $T = 20$.




(a) $\text{Var}\{\varepsilon\} = 0.25^2$ (b) $\text{Var}\{\varepsilon\} = 0.5^2$ (c) $\text{Var}\{\varepsilon\} = 1$

Figure 2.5: *The influence of noise on the performance of the methods.* The conditions of this experiment are exactly the same as those in figure 2.4, except that we added no noise. Observe how noise only acts as a limit to the attainable accuracy (at 10^{-3} in figure 2.4), without any further effects.

6 Conclusions and perspectives

We reviewed the theory of irregular sampling and the available algorithms and discussed its application to satellite imaging. Unfortunately none of the results on existence of a stable reconstruction formula do exactly apply to our case. Kadec-type results are too restrictive in terms of the allowed amplitude, whereas Beurling-Landau-type results are too restrictive in terms of the spectral support of the image. However Kadec's results do not make use of the regularity of the image⁶, nor of the regularity (the small spectral support) of the perturbation. On the other hand the results of Beurling and Landau do exploit the regularity of the perturbation in some way (the density condition can be translated in the case of perturbed sampling in terms of a bound on $D\varepsilon$), but not the regularity of the image, and the fact that the sampling is perturbed (not just a general irregular sampling). The question of whether these particular properties could eventually be used to relax Kadec's or Beurling's results is to the best of our knowledge still open.

Concerning reconstruction algorithms, none of the algorithms we analyzed provides a convergence result in the conditions we need to apply it. Nevertheless two of them, and the new pseudo-inverse algorithm we propose here, do provide reasonable approximations, below the noise level that usually occurs in satellite imagery.

It seems that our problem requires a shift in the formulation of the problem. We do no longer look for the conditions under which exact stable reconstruction is possible (this may not be possible, even if we take into account the regularity of the image and the perturbation), but for conditions under which the problem of reconstructing the regular samples *up to a certain accuracy* is reasonably well posed, and for criteria to determine which is the best possible *attainable accuracy* given the characteristics of the perturbation.

From an engineering point of view the choice between the three methods we outlined here depends on many factors, ranging from the importance of speed vs. accuracy, to the acquisition conditions. Under certain conditions, the pseudo-inverse produces results with the same accuracy as the ACT method, with less than 5 times less computations. In other conditions, the accuracy of the pseudo-inverse may be 10 times poorer than that of the ACT method.

We did not address here possible aliasing problems (or its equivalent in the irreg-

⁶in the sense that the image is not only band-limited but its Fourier coefficients also are small close to the Nyquist frequency

ular sampling case). The function was assumed to be perfectly band-limited. The instrument's transfer function was used only within $[-\pi, \pi]^2$ in order to simulate the kind of decay rate that we get from the sampling instrument. But beyond $[-\pi, \pi]^2$ either f or the transfer function were assumed to be zero, which is not strictly true.

A more accurate approach would be to assume f inside $V^p(H)$ for the known transfer function H . But computing orthogonal projections to $V^p(H)$ can be computationally very expensive. Alternatively we can use Aldroubi's algorithm to find the best approximation $\tilde{f} \in V^p(\phi)$ to the real f (given the irregular samples $f(\Lambda)$) by using a convenient, e.g. spline-like, generator ϕ which allows for fast computation, and at the same time is close to H . Note that in such a case the correct projector P in Aldroubi's algorithm should be a projection from L^p into $V^p(\phi)$ which is as close as possible to $V^p(H)$. This can be done by the oblique projection $P_{V^p(\phi) \perp V^p(H)}$ as suggested by Aldroubi and Feichtinger [2], Unser [171]. The only difference from the computational point of view is the calculation of the prefilter, which can be done with the same cost in the Fourier domain.

Appendix

A Proofs of convergence analysis.

In the proof of propositions 4 and 5 we use several times the following two results

Lemma 6. *If a function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ has spectral support $\text{supp}(\hat{f}) \subseteq R$ in a rectangular region $R = [-r_1, r_1] \times [-r_2, r_2]$, then*

$$\|Df\|_{L^2}^2 \leq (r_1^2 + r_2^2) \|f\|_{L^2}^2 \quad (2.61)$$

and therefore $Df = O(f)$.

This result follows simply from Parseval's formula and the fact that $\mathcal{F}(Df)(w) = w\hat{f}$. We shall apply it both to bound the norm of the image f and the perturbation ε .

The next result provides a convenient development of the inverse warp ϕ_ε^{-1} as a truncated power series in ε . (Recall that ϕ_ε^{-1} is the inverse of the warp $\phi_\varepsilon(x) = x + \varepsilon(x)$ mapping the regular grid \mathbb{Z}^2 to the irregular grid Λ).

Lemma 7. *The inverse warp ϕ_ε^{-1} can be expanded as*

$$\phi_\varepsilon^{-1}(x) = x - \varepsilon_1(x) + O(\varepsilon^3) \quad \text{where} \quad \varepsilon_1(x) = \varepsilon(x - \varepsilon(x)) \quad (2.62)$$

Proof. Let's consider the approximation $\tilde{\phi}_\varepsilon^{-1} = x - \varepsilon_1(x)$ and test how it approximates the inverse of $\phi_\varepsilon(x) = x + \varepsilon(x)$:

$$\begin{aligned} \tilde{\phi}_\varepsilon^{-1}(\phi_\varepsilon(x)) &= x + \varepsilon(x) - \varepsilon(x + \varepsilon(x)) - \varepsilon(x + \varepsilon(x)) \\ &= x - D\varepsilon(x) \cdot d(x) - D^2\varepsilon(d(x), d(x)) + O(\varepsilon^3). \end{aligned} \quad (2.63)$$

In the last equation we used the abbreviation $d(x) = \varepsilon(x) - \varepsilon(x + \varepsilon(x))$, which can itself be expanded as

$$d(x) = D\varepsilon(x) \cdot \varepsilon(x) + O(\varepsilon^2). \quad (2.64)$$

Using lemma 6 we conclude that $d(x) = O(\varepsilon^2)$, $D^2\varepsilon = O(\varepsilon)$ so that the second order term in equation (2.63) is actually $O(\varepsilon^5)$ and is absorbed in $O(\varepsilon^3)$. Furthermore the first order term $D\varepsilon(x) \cdot d(x)$ is actually $O(\varepsilon^3)$, because $D\varepsilon = O(\varepsilon)$ using once more lemma 6. So equation (2.63) can be rewritten as

$$\tilde{\phi}_\varepsilon^{-1}(\phi_\varepsilon(x)) = x + O(\varepsilon^3) = \phi_\varepsilon^{-1}(\phi_\varepsilon(x)) + O(\varepsilon^3). \quad (2.65)$$

□

Proof of proposition 4. We start by writing step by step the application of the operator $I - \alpha = A^-A^+$ to an image $s = \Delta_\Gamma f$ such that $\text{supp}(\hat{f}) \subseteq R$. (Recall that the sampling grid is $\Gamma = \mathbb{Z}^2$ and $R = [-\pi, \pi]^2$ is the corresponding reciprocal cell, so that the Shannon interpolation filter is $\mathcal{F}^{-1}(\mathbb{1}_R) = \text{sinc}$. We shall note by $\text{sinc}^c = \mathcal{F}^{-1}(\mathbb{1}_{R^c})$ the complementary filter. Note that $\text{sinc}^c * f = f - \text{sinc} * f$).

$$f = \text{sinc} * s \quad s = \Delta_\Gamma f \quad (2.66)$$

$$\tilde{f} = f \circ \phi_\varepsilon \quad \tilde{s} = \Delta_\Gamma \cdot \tilde{f} = A^+ s \quad (2.67)$$

Since the spectral support of \tilde{f} goes beyond R , it cannot be recovered from \tilde{s} . However, the interpolation of \tilde{s} yields a certain function \tilde{f}' with spectral support in R , that coincides with \tilde{f} at the sampling points

$$\tilde{f}' = \text{sinc} * \tilde{s} = \text{sinc} * \tilde{f} + \text{alias}(\tilde{f}) \quad \tilde{s} = \Delta_\Gamma \cdot \tilde{f}' \quad (2.68)$$

Next we apply $A^- \tilde{s} = \Delta_\Gamma \cdot ((\text{sinc} * \tilde{s}) \circ \phi_\varepsilon^{-1})$ in three steps: convolution (we just did it as $\tilde{f}' = \text{sinc} * \tilde{s}$), inverse warp, and sampling

$$f'_1 = \tilde{f}' \circ \phi_\varepsilon^{-1} \quad s_1 = \Delta_\Gamma \cdot f'_1 = A^- \tilde{s} \quad (2.69)$$

Finally consider the interpolation of s_1 :

$$f_1 = \text{sinc} * s_1 = \text{sinc} * f'_1 + \text{alias}(f'_1) \quad s_1 = \Delta_\Gamma \cdot f_1 = A^- \tilde{s} \quad (2.70)$$

In the sequel we shall use the abbreviation

$$S(f) = \text{sinc} * (\Delta_\Gamma \cdot f) = \text{sinc} * f + \text{alias}_R(f) \quad (2.71)$$

for the sampling-interpolation pair of operators (of course S is the identity if $\text{supp}(\hat{f}) \subseteq R$). With this notation the last equation becomes $f_1 = S(f'_1)$.

Now we are ready to compute the Taylor development of $\alpha f = f_1 - f$ with respect to ε . To do so we start by computing the third order Taylor development of \tilde{f} and \tilde{f}' .

$$\tilde{f} = f \quad + Df \cdot \varepsilon \quad + \frac{1}{2} D^2 f(\varepsilon, \varepsilon) + O(\varepsilon^3) \quad (2.72)$$

$$\begin{aligned} \tilde{f}' = f' & \quad + \text{sinc} * (Df \cdot \varepsilon) \quad + \frac{1}{2} \text{sinc} * (D^2 f(\varepsilon, \varepsilon)) \\ & \quad + \text{alias}_R(Df \cdot \varepsilon) \quad + \frac{1}{2} \text{alias}_R(D^2 f(\varepsilon, \varepsilon)) + O(\varepsilon^3) \end{aligned} \quad (2.73)$$

Since according lemma 7, the inverse warp is $\phi_\varepsilon^{-1} = x - \varepsilon_1(x) + O(\varepsilon^3)$, the Taylor development of f'_1 can be written

$$f'_1 = \tilde{f}' \quad - D\tilde{f}' \cdot \varepsilon_1 \quad + \frac{1}{2} D^2 f(\varepsilon_1, \varepsilon_1) + O(\varepsilon^3). \quad (2.74)$$

In order to express this development in terms of f and ε , we first expand $\varepsilon_1(x)$ as

$$\varepsilon_1(x) = \varepsilon(x - \varepsilon(x)) = \varepsilon(x) - D\varepsilon(x) \cdot \varepsilon(x) + O(\varepsilon^3) \quad (2.75)$$

$$= \varepsilon(x) + O(\varepsilon^2) \quad (2.76)$$

and secondly, we expand $D\tilde{f}'$ and $D^2\tilde{f}'$ up to the right order using equation (2.73)

$$\begin{aligned} D\tilde{f}' = Df & \quad + \text{sinc} * (D^2 f(\varepsilon, \cdot) + Df \cdot D\varepsilon) \quad + \\ & \quad + D \text{alias}_R(Df \cdot \varepsilon) \quad + O(\varepsilon^2) \end{aligned} \quad (2.77)$$

$$D^2\tilde{f}' = D^2 f + O(\varepsilon). \quad (2.78)$$

Substituting the last three equations in equation (2.74) we obtain

$$\begin{aligned} f'_1 = f & \quad + \text{sinc} * (Df \cdot \varepsilon) \quad + \text{sinc} * (D^2 f(\varepsilon, \varepsilon)) \\ & \quad + \text{alias}_R(Df \cdot \varepsilon) \quad + \text{alias}_R(D^2 f(\varepsilon, \varepsilon)) \quad \left. \vphantom{f'_1} \right\} = \tilde{f}' + O(\varepsilon^3) \\ & \quad - Df \cdot \varepsilon_1 - (\text{sinc} * (D^2 f(\varepsilon, \cdot) + Df \cdot D\varepsilon)) \cdot \varepsilon_1 \\ & \quad \quad - (D \text{alias}_R(Df \cdot \varepsilon)) \cdot \varepsilon_1 \quad \left. \vphantom{f'_1} \right\} = D\tilde{f}' \cdot \varepsilon_1 + O(\varepsilon^3) \\ & \quad + D^2 f(\varepsilon_1, \varepsilon_1) \quad \left. \vphantom{f'_1} \right\} = D^2\tilde{f}'(\varepsilon_1, \varepsilon_1) + O(\varepsilon^3) \\ & \quad + O(\varepsilon^3). \end{aligned} \quad (2.79)$$

Sampling and reinterpolation of f'_1 introduces new aliases in all terms except for \tilde{f}' whose spectral support is already in R :

$$\begin{aligned}
f_1 = f & \left. \begin{array}{l} \boxed{+\text{sinc} * (Df \cdot \varepsilon)} \quad \boxed{+\frac{1}{2}\text{sinc} * (D^2f(\varepsilon, \varepsilon))} \\ \boxed{+\text{alias}_R(Df \cdot \varepsilon)} \quad \boxed{+\frac{1}{2}\text{alias}_R(D^2f(\varepsilon, \varepsilon))} \end{array} \right\} = \tilde{f}' + O(\varepsilon^3) \\
& \left. \begin{array}{l} \boxed{-\text{sinc} * (Df \cdot \varepsilon_1)} \quad \boxed{-\text{sinc} * ((\text{sinc} * (D^2f(\varepsilon, \cdot)) \cdot \varepsilon_1)} \quad \boxed{-\text{sinc} * ((\text{sinc} * (Df \cdot D\varepsilon)) \cdot \varepsilon_1)} \\ \boxed{-\text{alias}_R(Df \cdot \varepsilon_1)} \quad \boxed{-\text{alias}_R((\text{sinc} * (D^2f(\varepsilon, \cdot)) \cdot \varepsilon_1)} \quad \boxed{-\text{alias}_R((\text{sinc} * (Df \cdot D\varepsilon)) \cdot \varepsilon_1)} \\ \quad \quad \quad -\text{sinc} * (D \text{alias}_R(Df \cdot \varepsilon)) \cdot \varepsilon_1 \\ \quad \quad \quad -\text{alias}_R(D \text{alias}_R(Df \cdot \varepsilon)) \cdot \varepsilon_1 \end{array} \right\} = \begin{array}{l} S(D\tilde{f}' \cdot \varepsilon_1) \\ +O(\varepsilon^3) \end{array} \\
& \left. \begin{array}{l} \boxed{+\frac{1}{2}\text{sinc} * (D^2f(\varepsilon_1, \varepsilon_1))} \\ \boxed{+\frac{1}{2}\text{alias}_R(D^2f(\varepsilon_1, \varepsilon_1))} \end{array} \right\} = S(D^2\tilde{f}'(\varepsilon_1, \varepsilon_1)) + O(\varepsilon^3) \\
& + O(\varepsilon^3). \tag{2.80}
\end{aligned}$$

Finally we use the expression of ε_1 given in equation (2.75) in order to group the terms

$$\begin{aligned}
\boxed{\square} &= \text{sinc} * [(Df \cdot \varepsilon) - (Df \cdot (\varepsilon - D\varepsilon \cdot \varepsilon + O(\varepsilon^3)))] - ((\text{sinc} * (Df \cdot D\varepsilon)) \cdot (\varepsilon + O(\varepsilon^2))) \\
&= \text{sinc} * (\text{sinc}^c * [Df \cdot D\varepsilon] \cdot \varepsilon) + O(\varepsilon^3) \tag{2.81}
\end{aligned}$$

$$\begin{aligned}
\boxed{\square} &= \text{sinc} * \left[\frac{1}{2}D^2f(\varepsilon, \varepsilon) + \frac{1}{2}D^2f(\varepsilon + O(\varepsilon^2), \varepsilon + O(\varepsilon^2)) - ((\text{sinc} * D^2f(\varepsilon, \cdot)) \cdot (\varepsilon + O(\varepsilon^2))) \right] \\
&= \text{sinc} * (\text{sinc}^c * [D^2f(\varepsilon, \cdot)] \cdot \varepsilon) + O(\varepsilon^3) \tag{2.82}
\end{aligned}$$

Similarly we conclude that

$$\boxed{\square\square} = \text{alias}_R(\text{sinc}^c * [Df \cdot D\varepsilon] \cdot \varepsilon) + O(\varepsilon^3) \tag{2.83}$$

$$\boxed{\square} = \text{alias}_R(\text{sinc}^c * [D^2f(\varepsilon, \cdot)] \cdot \varepsilon) + O(\varepsilon^3) \tag{2.84}$$

Taking into account that $D(Df \cdot \varepsilon) = D^2f(\varepsilon, \cdot) + Df \cdot D\varepsilon$, and recalling the notation $S(g) = \text{sinc}^c * g + \text{alias}_R(g)$ for the sampling-interpolation pair, we can summarize equation (2.80) as follows:

$$f'_1 = f + S([\text{sinc}^c * (D(Df \cdot \varepsilon)) - D \text{alias}_R(Df \cdot \varepsilon)] \cdot \varepsilon) + O(\varepsilon^3) \tag{2.85}$$

Then, equation (2.41) follows by sampling once more (since sampling after applying S is equivalent to just sampling). \square

Proof of proposition 5. We abbreviate the main term in equation (2.85) as

$$C_\varepsilon f = S([\text{sinc}^c * (D(Df \cdot \varepsilon)) - D \text{alias}_R(Df \cdot \varepsilon)] \cdot \varepsilon) \quad (2.86)$$

and we further decompose both components of the vector $C_\varepsilon = \begin{pmatrix} C_{1,\varepsilon} \\ C_{2,\varepsilon} \end{pmatrix}$, as well as $\varepsilon = \begin{pmatrix} \varepsilon_1 \\ \varepsilon_2 \end{pmatrix}$ and $D = \begin{pmatrix} D_1 \\ D_2 \end{pmatrix}$:

$$C_{j,\varepsilon} f = S(\text{sinc}^c * (D_j(D_1 f \varepsilon_1 + D_2 f \varepsilon_2)) - D_j \text{alias}_D^*(D_1 f \varepsilon_1 + D_2 f \varepsilon_2)) + o(\varepsilon^3) \quad \text{for } j = 1, 2 \quad (2.87)$$

This way αf can be written as

$$\alpha f = C_\varepsilon f \cdot \varepsilon + o(\varepsilon^3) \quad (2.88)$$

$$= C_{1,\varepsilon} f \varepsilon_1 + C_{2,\varepsilon} f \varepsilon_2 + o(\varepsilon^3) \quad (2.89)$$

To find an upper bound of $\|\alpha f\|_{L^2}$ we apply first the triangular inequality to the scalar product $D_\varepsilon f \cdot \varepsilon$, and then we compute the max of ε_j over the whole domain:

$$\|\alpha f\|_2 \leq \|C_{1,\varepsilon} f \varepsilon_1\|_2 + \|C_{2,\varepsilon} f \varepsilon_2\|_2 + o(\varepsilon^3) \quad (2.90)$$

$$\leq \|C_{1,\varepsilon} f\|_2 \|\varepsilon_1\|_\infty + \|C_{2,\varepsilon} f\|_2 \|\varepsilon_2\|_\infty + o(\varepsilon^3) \quad (2.91)$$

$$(2.92)$$

(Observe that we could as well have applied Cauchy-Schwartz, but this would introduce a factor in $\|\varepsilon\|_{L^2}$ which depends on the image size). Next, we apply once again the triangular inequality, this time to $\|C_{j,\varepsilon} f\|$:

$$\begin{aligned} \|C_{j,\varepsilon} f\| &\leq \|\text{sinc}^c * D_j(D_1 f \varepsilon_1)\| && + \|D_j \text{alias}_R(D_1 f \varepsilon_1)\| && + \\ &\|\text{sinc}^c * D_j(D_2 f \varepsilon_2)\| && + \|D_j \text{alias}_D^*(D_2 f \varepsilon_2)\| && + o(\varepsilon^3) \end{aligned} \quad (2.93)$$

Finally, each of the four terms (for $j = 1$) can be upper bounded, by taking into account the link between differentiation in the Fourier domain, the aliases and the different spectral supports:

$$\|\text{sinc}^c * D_1(D_1 f \varepsilon_1)\| \leq \sum_{k=1}^N |A_{1,k}| (\|\hat{f} \mathbb{1}_{B_k^+}\| + \|\hat{f} \mathbb{1}_{B_k^-}\|) \pi^2 \quad (2.94)$$

$$\|\text{sinc}^c * D_1(D_2 f \varepsilon_2)\| \leq \sum_{k=1}^N |A_{2,k}| (\|\hat{f} \mathbb{1}_{B_k^+}\| + \|\hat{f} \mathbb{1}_{B_k^-}\|) \pi(\pi + |w_{1,k}|) \quad (2.95)$$

$$\|D_1 \text{alias}_D^*(D_1 f \varepsilon_1)\| \leq \sum_{k=1}^N |A_{1,k}| (\|\hat{f} \mathbb{1}_{B_k^+}\| + \|\hat{f} \mathbb{1}_{B_k^-}\|) \left(\pi - \frac{1}{2}|w_{1,k}|\right)^2 \quad (2.96)$$

$$\|D_1 \text{alias}_D^*(D_2 f \varepsilon_2)\| \leq \sum_{k=1}^N |A_{2,k}| (\|\hat{f} \mathbb{1}_{B_k^+}\| + \|\hat{f} \mathbb{1}_{B_k^-}\|) \pi^2 \quad (2.97)$$

A similar upper bound can be found for $j = 2$, and by substituting $|w_{1,k}|$ and $|w_{2,k}|$ by the common upper bound $|w_k|_\infty$ (the largest coordinate in module), we obtain (2.44), after simply rearranging the terms. This completes the proof of the first part of the proposition.

To show the second part we just observe that the upper bound $\frac{\text{area}(B_k^+)}{\text{area}([- \pi, \pi]^2)} \|\hat{f}\|$ for $\|\hat{f}\|_{B_k^+}$, can be itself upper-bounded by:

$$\frac{\text{area}(B_k^+)}{\text{area}([- \pi, \pi]^2)} = (\lambda_1 + \lambda_2) - \lambda_1 \lambda_2 \leq 2 \sup\{\lambda_1, \lambda_2\} = \frac{2|w_k|_\infty}{\pi} \quad (2.98)$$

where $\lambda_j = \frac{|w_{j,k}|}{2\pi}$. □

Sample application As an application of the upper bound provided by proposition 5, assume that ε has N vibrating modes, with amplitude $|A_{j,k}| < A$ and frequency $|w_k|_\infty < \frac{\pi}{T}$. Then, $\|\varepsilon_j\|_\infty < NA$, so

$$\frac{\|\alpha f\|}{\|f\|} \leq (4NA)^2 4\pi^2 \frac{2}{T} \left(1 + \frac{1}{2T}\right)$$

To have convergence in the pseudo-inverse algorithm, this value should be smaller than 1 (minus the term in $O(\varepsilon^3)$), *i.e.*:

$$A \leq \frac{1}{8N\pi \sqrt{\frac{2}{T} \left(1 + \frac{1}{2T}\right)}}$$

For instance, for $T = 10$ we may obtain convergence for $A \leq 0.18 \frac{1}{N}$. For a single vibrating mode ($N = 1$) this means that the maximal perturbation amplitude $2A$ of the sinusoidal functions shouldn't exceed 0.36 pixels, which is already larger than Kadec's result. However, since we do not control the third order term, this does not ensure convergence.

Perspectives. The upper bound for $\|\alpha s\|$ we just showed is not very tight, and can be improved in several ways.

First the regularity hypothesis we assume on f (white noise) is very weak. Natural images are usually much more regular, with a faster decay of Fourier coefficients. If we take into account this decay and the decay imposed by the transfer function, we can obtain a tighter upper bound.

Secondly, our upper bound considers the worst case, which is rarely attained, especially when ε contains many Fourier coefficients. If we consider mean errors (by taking expected values instead of sup's, the upper bounds would be still smaller. For instance

the expected value of $\|\varepsilon_j\|_\infty$ for ε_j with N Fourier coefficients having amplitude $\leq A$, but a random phase, is $\sqrt{N}A$, which is considerably smaller than the preceding upper-bound NA .

Finally, this upper bound should be applied locally as explained in the main text. Since this allows to reduce the amplitude (to the local variations, not taking into account very low-frequency components of the perturbation), as well as the number N of terms defining ε , since only a few are significant in a small domain.

Part II

Non-Shannon Interpolation

Chapter 3

Interpolation of Digital Terrain Models using AMLE and Related Methods *

Abstract

Interpolation of Digital Elevation Models (DEMs) becomes necessary in many situations. For instance when constructing them from contour lines (available *e.g.* from non-digital cartography), or from disparity maps based on pairs of stereoscopic views, which often leaves large areas where point correspondences cannot be found reliably.

The Absolutely Minimizing Lipschitz Extension (AMLE) model is singled out as the simplest interpolation method satisfying a set of natural requirements. In particular, a maximum principle is proven, which guarantees not to introduce unnatural oscillations which is a major problem with many classical methods. We then discuss the links between the AMLE and other existing methods. In particular we show its relation with geodesic distance transformation. We also relate the AMLE to the thin plate method, that can be obtained by a prolongation of the axiomatic arguments leading to the AMLE, and addresses the major disadvantage of the AMLE model, namely its inability to interpolate slopes as it does for values. Nevertheless, in order to interpolate slopes, we have to give up the maximum principle and authorize the appearance of oscillations. We also discuss the possible link between the AMLE and Kriging methods that are the most widely used in the geoscience literature.

We end by numerical comparisons of the different methods. Our experiments show that despite its inability to interpolate slopes, the AMLE produces qualitatively better results than other methods, since it better preserves ridges. Hence we propose a simple combination of the AMLE and Kriging models which preserves the slope-interpolating and ridge-preserving capabilities of Kriging and AMLE respectively. Either the combined method or the AMLE showed the best approximation results in all cases both in qualitative and quantitative terms.

*A preliminary version of this chapter was jointly published with Frédéric Cao, Yann Gousseau and Bernard Rougé in IEEE Trans. GARS [8]. Early versions of this work were submitted to CNES as an internal report [9], and presented at the European Workshop of Mathematical Morphology [10]. Another application of thin-plate splines to elastic fingerprint matching was published with Laurent Cohen in the proceedings of the IEEE WACV [11].

1 Introduction

In this chapter, we study the problem of interpolating two dimensional data. This is motivated by several problems that may be encountered when trying to reconstruct digital elevation models:

- For instance, for maps which are not yet available in digital form, the elevation is only known on a limited number of level curves, and we then have to interpolate in order to retrieve a three dimensional model of elevation.
- One way to construct elevation models is to find points of correspondence between two images from a stereo vision pair of a scene. Most existing matching algorithms fail on some parts of the image and no elevation can be retrieved [57]. Interpolation is then also necessary.

In both cases, we may assume that the elevation is regular in the areas where it is *a priori* unknown. In the first case, the lines whose level is known are in general chosen such that they are representative of the real elevation. Therefore, the elevation between two adjacent lines always lies between the corresponding levels. Moreover it is not likely that oscillations between both lines can be significant, else some additional level lines would have also been represented. This implies that a reasonable interpolation method must not artificially create some oscillations between given level curves, otherwise it may create some arbitrary information which cannot be inferred from the data. In the second case, correlation is commonly used to find point correspondences. In this case, one of the major reasons why matching fails is that the variations of the image are not high enough in some region, which implies that the corresponding elevation is likely to be regular.

An interpolation model called AMLE was introduced and characterized by Aronsson [24, 25]. It was later proposed for digital images by Caselles, Morel, and Sbert [42] as the most simple interpolator satisfying a set of natural axioms. We think that these axioms are also very well suited for the interpolation of digital elevation models. We recall these axioms in section 2. A very nice property of the AMLE, that can be mathematically proved, is that it does not create any artifacts nor oscillations. This model is a real interpolation model that exactly fits the data¹. It can interpolate values on isolated curves and even isolated points (for instance mountain or hill top).

¹Here we mean a model that actually fits the data on curves and isolated points, and doesn't have zero capacity for points like Laplace's operator for instance (see section 2 for a discussion).

We show in section 3 how the AMLE can be related to other methods already used for DEM interpolation. First, it can be viewed as the stationary state of iterated geodesic distance transformations (see section 3.1). Second, it may happen that the elevation is known not only on isolated points but on a domain and we want to reconstruct the elevation outside this set. In this case, we can take advantage of the value on the boundary of the domain but also of some higher order information such as the slope of the data. The AMLE cannot handle more than zero-th order information and neither do geodesic distance interpolators. The thin plate model discussed in section 3.2 can be viewed as an attempt to generalize the axiomatic approach leading to the AMLE. It allows to interpolate values as well as gradient fields on the boundary of the interpolation domain, but it does not guarantee to avoid oscillations as the AMLE model does. Kriging methods (see section 3.3) are also related to the thin-plate model in some cases, and we compare the AMLE, the thin-plate and Kriging in section 4 by showing some interpolation results for both methods. The experiments suggested a simple combination of the AMLE and Kriging models which outperformed all the other methods in most cases. Finally in section 5 we discuss some issues for future research.

2 The AMLE model

In this section, we introduce the Absolutely Minimizing Lipschitz Extension (AMLE) model. It has been introduced by Caselles, Morel and Sbert in [42] as one of the pertinent models for image interpolation, though it had been previously studied from a completely theoretical point of view by Aronsson [24, 25] and Jensen [102]. The approach they use is axiomatic, that is, from a small number of natural postulates, they derive a classification of interpolation operators. Among the ones satisfying the largest number of these natural properties, the AMLE model will be particularly relevant in the problems we are interested in. In particular, it will be very efficient to interpolate data between level lines. This shall be interesting for elevation reconstruction from a scanned map where only a few level lines (iso-level curves) are available. We shall also use it to interpolate scattered data obtained from a correlation between stereo pairs of a scene.

Let us now give the main ideas leading to the AMLE model. We assume that the elevation we want to interpolate is the interior of a domain Ω in the plane. We assume that we only know the value φ of elevation on $\partial\Omega$, the boundary of Ω . Note that this is re-

alistic in the interpolation of iso-level lines. We assume that it is possible to interpolate φ inside Ω . Let us denote by $E(\Omega, \varphi) : \Omega \rightarrow \mathbb{R}$ the obtained elevation. We now review some natural properties that $E(\Omega, \varphi)$ should satisfy.

(P1) The interpolation is stable, that is to say, if we interpolate the values of $E(\Omega, \varphi)$ on a subdomain of Ω , the result does not change. More precisely, for any subdomain $\Omega' \subset \Omega$, we have

$$E(\Omega', E(\Omega, \varphi)|_{\partial\Omega'}) = E(\Omega, \varphi)|_{\Omega'}, \quad (3.1)$$

where the subscript $|A$ designs the restriction of a function to the set A . This property was introduced by [24] and lead to his geometric characterization of AMLE.

(P2) The interpolation respects the global elevation ordering, that is

$$\text{if } \varphi \leq \psi \text{ on } \partial\Omega, \text{ then } E(\Omega, \varphi) \leq E(\Omega, \psi). \quad (3.2)$$

(P3) Finally, we give a more technical property giving the behavior of E on second order polynomials. It consists in assuming the existence of a function F such that if $u(x) = c + \langle p, x - x_0 \rangle + \frac{1}{2} \langle A(x - x_0), (x - x_0) \rangle$ is a second order polynomial with coefficients given by the constant c , the real 2-vector p and the real 2×2 matrix A , then

$$\lim_{r \rightarrow 0} \frac{E(D(x_0, r), u|_{D(x_0, r)})(x_0) - u(x_0)}{r^2/2} = F(A, p, c, x_0). \quad (3.3)$$

To these three basic axioms, we add a set of natural geometric properties.

(P4) The interpolation does not depend on the position of Ω . In other terms, it commutes with translation.

(P5) It does not depend on the orientation of Ω either. Put another way, it commutes with plane rotations.

(P6) It is also scale invariant (*i.e.* commutes with dilations).

(P7) The reference level is arbitrary, hence the interpolation must commute with the addition of a fixed value.

(P8) The elevation unit is also arbitrary. Therefore, the interpolation has to commute with multiplication by a fixed constant.

Before we give the general form of an interpolation operator satisfying these properties, let us give some notations.

If $u : \Omega \subset \mathbb{R}^2 \rightarrow \mathbb{R}$ is a real-valued function taking values in a region Ω of the plane, the gradient of u denoted by Du , is the vector with coordinates (u_x, u_y) of the partial derivatives of u with respect to x and y . We denote by Du^\perp , the vector obtained from Du by a rotation of angle $\pi/2$ (the orientation does not matter; we arbitrarily chose it such that (Du, Du^\perp) is counterclockwise if $Du \neq 0$, *e.g.* we can choose the convention

$Du^\perp = (-u_y, u_x)$). We also denote by D^2u the Hessian matrix of u , that is to say, the symmetric matrix whose coefficients are the second order derivatives of u . As usual, we canonically associate a quadratic form to D^2u . Using this quadratic form we can compute the second order derivative in any direction as well as mixed derivatives. In particular we use the notation $u_{\xi\xi} = D^2u \left(\frac{Du}{|Du|}, \frac{Du}{|Du|} \right)$ to represent the second derivative of u in the direction of the gradient, and similarly $u_{\eta\eta} = D^2u \left(\frac{Du^\perp}{|Du^\perp|}, \frac{Du^\perp}{|Du^\perp|} \right)$ and $u_{\eta\xi} = D^2u \left(\frac{Du^\perp}{|Du^\perp|}, \frac{Du}{|Du|} \right)$ to represent the second derivative of u in the direction of the level curves, and the mixed derivative, respectively. The following result characterizes all interpolators satisfying the required axioms:

Theorem 10 (Caselles, Morel Sbert [42]). *Assume that $E(\Omega, \varphi)$ satisfies the properties (P1)-(P8). Then, it is a viscosity solution of an equation of the form*

$$G(u_{\xi\xi}, u_{\eta\eta}, u_{\eta\xi}) = 0, \quad (3.4)$$

where $G(a, b, c)$ is positively homogeneous of degree 1 and nondecreasing with respect to the matrix $\begin{pmatrix} a & c \\ c & b \end{pmatrix}$.

If we assume that G is differentiable at the point $(0, 0, 0)$, then, G is linear, that is $G(a, b, c) = \alpha a + \beta b + \gamma c$ where (α, β, γ) satisfies the property $\alpha\beta - \gamma^2 \geq 0$.

The viscosity solution theory is the correct mathematical setting for nonlinear elliptic equations of the type 3.4. We refer the reader to [52, 65] for complete details.

In the following, we shall assume that the function G in Equation (3.4) is differentiable, and thus linear. We see, that the most simple equations correspond to the following three cases: ²

1. $G(a, b, c) = a + b$. In this case, the equation is $\Delta u = 0$, that is Laplace equation.
2. $G(a, b, c) = a$, yielding the equation $D^2u \left(\frac{Du}{|Du|}, \frac{Du}{|Du|} \right) = 0$, which is Aronsson's AMLE model.
3. $G(a, b, c) = b$, yielding the equation $D^2u \left(\frac{Du^\perp}{|Du^\perp|}, \frac{Du^\perp}{|Du^\perp|} \right) = 0$. This equation means that the curvature of the iso-level lines equals 0. As a consequence, these curves are straight lines.

² The other cases we could derive from theorem 10 are also excluded if we want an interpolator satisfying (P1)-(P8), which interpolate both curves and points and lead to regular interpolators (finite $\|Du\|_{L^\infty(\Omega)}$ if the boundary data is regular too (Lipschitz)). A discussion in [42] shows this on a simple example where the boundary data are available on a circle and its center.

Laplace's equation is well known for regular interpolation since it can describe the equilibrium temperature distribution in a domain with source on the boundary. Nevertheless, by classical results of potential theory, this model does not allow to fix values at isolated points, and we say that it has *zero capacity for points*. For instance, if $\Omega = D(0, 1) \setminus \{0\}$ in the plane (the pointed unit disk) and if we fix $\varphi = 0$ on the unit circle and $\varphi = 1$ at the origin, the solution of Laplace equation is identically zero, letting the value at the origin ignored! It is possible to prove that any linear combination of the form $\alpha D^2 u \left(\frac{Du}{|Du|}, \frac{Du}{|Du|} \right) + \beta D^2 u \left(\frac{Du^\perp}{|Du|}, \frac{Du^\perp}{|Du|} \right)$ with positive α and β satisfies the same undesirable property.

The equation $D^2 u \left(\frac{Du^\perp}{|Du|}, \frac{Du^\perp}{|Du|} \right) = 0$ may not have a unique solution. In [123], Masnou and Morel found a way to choose the best solution by introducing a variational condition. This algorithm, called disocclusion gives impressive results in restoring images, but the domain Ω must be simply connected (which is a too strong condition for our application) and the fact that iso-level curves are straight curves somehow gives unrealistic elevation models.

On the contrary, the AMLE model can interpolate isolated values. In the case of the pointed disk, the solution is the one we can expect, that is $u(x) = 1 - |x|$. The AMLE model is mathematically well posed as exposed in the following theorem.

Theorem 11 (Aronsson [24] and Jensen [102]). *Let Ω be a bounded domain. Assume that φ is continuous on $\partial\Omega$. Then, there exists a unique viscosity solution of*

$$D^2 u \left(\frac{Du}{|Du|}, \frac{Du}{|Du|} \right) = 0 \quad (3.5)$$

with boundary value equal to φ . Moreover, the AMLE satisfies the properties (P1)-(P8).

(A generalization of this result, which relaxes the continuity assumption on the boundary data φ , was proven by Cao in [41]).

The reason why this model is named AMLE is given by the following proposition.

Theorem 12 (Jensen [102]). *Let Ω be a bounded domain and φ be a Lipschitz function on $\partial\Omega$.³ Then the AMLE is the unique function u interpolating φ and satisfying the property*

$$\begin{aligned} \forall \Omega' \subset \Omega, \forall w \text{ Lipschitz in } \partial\Omega', w = u \text{ on } \partial\Omega' \Rightarrow \\ \|Du\|_\infty \leq \|Dw\|_\infty. \end{aligned}$$

³ Here this means that the sup $\frac{|\varphi(x) - \varphi(y)|}{|x - y|}$ with respect to $x, y \in \partial\Omega$ is finite, where $|x - y|$ is the usual euclidean distance in \mathbb{R}^2 , not the geodesic distance inside $\partial\Omega$.

This means that the *AMLE* minimizes the Lipschitz constant on any subdomain, which is an important property when interpolating elevation models because it is a way to ensure that the interpolation will not create information that is not present in the original data.

Many interpolation methods may produce oscillations (Gibbs effects, overshoots) which are not apparent from the original data. Such oscillations appear rather as a byproduct of regularity assumptions which do not match the given data. Theorem 12 ensures that the AMLE interpolator cannot create such oscillations.

Section 4 shows some experimental results using this model, and in appendix C.1 we give some details on the numerical solution of this equation and its computational complexity.

3 Relations to Previous Work on DEM interpolation

Whereas the AMLE model has been recently proposed for image processing in [42], it is essentially new in the geoscience domain [7]. In this section we explore the close relationships between AMLE and other classical methods that have been proposed in the geoscience literature.

3.1 Geodesic distance transformations.

Interpolators based on geodesic distances are to the best of our knowledge the main tool for constructing DEMs from contour lines, and fast algorithms have been proposed by Pierre Soille [163, 165]. Starting from the assumption that the boundary can be partitioned into two disjoint regions $\partial\Omega = \partial\Omega_1 \cup \partial\Omega_2$, such that the boundary conditions are constant in each of the regions, *i.e.* $u|_{\partial\Omega_1} = c_1 \in \mathbb{R}$ and $u|_{\partial\Omega_2} = c_2 \in \mathbb{R}$, the geodesic distance interpolator can be defined as:

$$u(x) = \frac{c_2 d_{\Omega,g}(x, \partial\Omega_1) + c_1 d_{\Omega,g}(x, \partial\Omega_2)}{d_{\Omega,g}(x, \partial\Omega_1) + d_{\Omega,g}(x, \partial\Omega_2)} \quad (3.6)$$

where $d_{\Omega,g}(x, A)$ is the generalized geodesic distance from a point x to a set A , *i.e.* the minimal usual geometric length of a path C entirely contained in Ω , which minimizes the geodesic distance $\int_C g$ among all paths joining x to a point $y \in A$. Whereas in [163] a constant weighting function g was used, in a later article [165] it was proposed to use the Euclidean distance to the border $g(x) = [d(x, \partial\Omega)]^c$ as a weighting function, in order to avoid certain artifacts occurring when geodesic paths C become tangent to the

border.⁴

This produces results which are very close to the AMLE model, which is not surprising since the AMLE model can under certain conditions be obtained as a fixed point of weighted geodesic distance interpolations. In fact, assuming that *e.g.* $c_1 \leq c_2$ consider the following iteration:

$$u_n(x) = \frac{c_2 L_1^n(x) + L_2^n(x) c_1}{L_1^n(x) + L_2^n(x)} \quad \text{where} \quad L_i^n(x) = \begin{cases} d_{\Omega,1}(x, \partial\Omega_i) & \text{if } n = 0 \\ d_{\Omega,|Du_{n-1}|}(x, \partial\Omega_i) & \text{if } n > 0 \end{cases}. \quad (3.7)$$

If the iteration converges to a fixed point

$$u_\infty(x) = \frac{c_2 L_1^\infty(x) + L_2^\infty(x) c_1}{L_1^\infty(x) + L_2^\infty(x)} \quad \text{with} \quad L_i^\infty(x) = d_{\Omega,|Du_\infty|}(x, \partial\Omega_i) \quad (3.8)$$

and this fixed point is twice differentiable then u_∞ is AMLE. Indeed, the geodesic paths associated to a C^1 potential $|Du_\infty|$ can be shown to be the gradient curves⁵ of u_∞ . Hence, equation (3.8) means that u_∞ is linear along its own gradient curves, *i.e.* $\|Du_\infty\|$ is constant in the direction of Du_∞ , or put another way u_∞ satisfies the AMLE equation $D^2u(Du, Du) = 0$. The details of the proof are developed in section A.2, and in section A.3 we extend this result under weaker assumptions on the fixed point. In general, however, the AMLE interpolation needn't satisfy even the weakest assumptions, so fixed points and AMLE interpolations are not always equivalent, but our simulations show that they are very close even when the hypothesis needed for equivalence are not satisfied.

This result has two important implications: (i) the different variations of (non-iterated) geodesic distance transform methods can be interpreted as truncated implementations of the AMLE model and (ii) they can be used to write faster algorithms to solve the AMLE in the particular case of contour line interpolation.

3.2 The thin plate model

A major drawback of the AMLE model is that it cannot interpolate slopes, it can only interpolate boundary data. As we shall see later in this section and in section 4 this does not always represent a serious problem, and there are many reasons for keeping the AMLE method despite this drawback. However, the fact that it cannot interpolate slopes may result in flat mountain tops and slope discontinuities across some level

⁴Here the complement operation is defined as $g^c = -g + \max_\Omega g$. It is essentially a negation with a shift in order to avoid the weight g^c from becoming negative or zero.

⁵A gradient curve of a function g is a curve such that its tangent vector is always parallel to Dg

curves. This is not really serious, at the sides of a mountain or when looking at the DEM as a gray-level image, but it becomes sometimes visible when representing the DEM as a 3D surface, producing undesirable artifacts.

Actually the fact that the AMLE cannot interpolate slopes comes from the axiomatic approach that leads to this model. Hence, in order to avoid these artifacts, we have to go back to our assumptions and do the necessary adaptations. First, we will require our operator E to interpolate not only boundary data, but also boundary slopes. Put another way, our operator will take the form $E(\Omega, \varphi, \psi) : \Omega \rightarrow \mathbb{R}$ where $\varphi = u|_{\partial\Omega}$ is the boundary data as usual, and $\psi = \frac{\partial u}{\partial n}|_{\partial\Omega}$ is the derivative of u in the direction of the inner normal n to Ω . The stability (P1) and invariance properties (P4-P8) can be trivially generalized to the new notation. The other two axioms will require, however, certain adaptations.

First, once we require E to interpolate slopes as well, (P2) does no longer make sense. A reasonable generalization is to require an order for the slopes as well as for the data:

(P2') If $\varphi_1 \leq \varphi_2$ and $\psi_1 \leq \psi_2$ on $\partial\Omega$ then $E(\Omega, \varphi_1, \psi_1) \leq E(\Omega, \varphi_2, \psi_2)$.

Similarly, since it is not possible to interpolate data and slopes with second order operators, we will require at least fourth order. Hence, we have to modify the regularity property as follows:

(P3') If $u(x)$ is a fourth order polynomial with coefficients c_0, \dots, c_m then

$$\lim_{r \rightarrow 0} \frac{E(u, u|_{D(x_0, r)}, \frac{\partial u}{\partial n}|_{D(x_0, r)})(x_0) - u(x_0)}{r^4} = F(c_m, \dots, c_0, x_0). \quad (3.9)$$

Among fourth order differential operators the biharmonic operator $\Delta^2 = \frac{\partial^4}{\partial x^4} + 2\frac{\partial^2}{\partial x^2 \partial y^2} + \frac{\partial^4}{\partial y^4}$ is well known, and its application to surface interpolation dates back to Duchon [66], Meinguet [127] and Franke [77]. The interpolating operator $u = E(\Omega, \varphi, \psi)$ consists of solving the biharmonic problem

$$\begin{cases} \Delta^2 u|_{\Omega} = 0 \\ u|_{\partial\Omega} = \varphi \\ \frac{\partial u}{\partial n}|_{\partial\Omega} = \psi \end{cases} \quad (3.10)$$

and is known to have a unique solution under quite general conditions on Ω , φ and ψ . This interpolator was called thin-plate spline because it closely models the shape taken by a thin metal plate when forced to the given boundary conditions. In such a situation, the plate minimizes its bending energy

$$J(u) = \iint_{\Omega} \left(\frac{\partial^2 u}{\partial x^2} \right)^2 + 2 \left(\frac{\partial^2 u}{\partial x \partial y} \right)^2 + \left(\frac{\partial^2 u}{\partial y^2} \right)^2 dx dy \quad (3.11)$$

subject to the boundary conditions $u|_{\partial\Omega} = \varphi$ and $\frac{\partial u}{\partial n}|_{\partial\Omega} = \psi$. We can easily show that the biharmonic equation (3.10) is the Euler-Lagrange of the bending energy (3.11).

This thin-plate spline model has been introduced as an interpolation tool for spatial data by Mitas and Mitasova [132] and the references therein. The authors also discuss some variations of the model like the thin-plate spline with tension with the aim of avoiding the oscillations that it may produce. This approach however relies on a trade-off between Laplacian interpolation (which has zero capacity for points) and biharmonic interpolation, and requires the empirical selection of a tension parameter. So in the rest of this section we shall concentrate on the thin-plate model.

Both from the PDE formulation and from the variational formulation, it is straightforward to verify that the thin-plate spline operator satisfies the invariance properties (P4-P8). Similarly, the stability property (P1') and the regularity property (P3') are direct consequences of its PDE formulation. The maximum principle (P2'), however, remains more subtle. Actually, Boggio and Hadamard [35, 92] conjectured that it should be true for quite general domains Ω , since Boggio showed that it was true for circular domains. Nevertheless 40 years later a number of counterexamples disproved the conjecture. For instance, Coffman and Duffin [47] showed it is false for rectangular Ω , and Garabedian [1, 79] showed it is false for elliptic domains with eccentricity two. In any case, the maximum principle remains valid for small perturbations of circular domains [89].

Unfortunately there is no equivalent for Theorem 12 either. We do know [137] that for bounded and Lipschitz domains Ω , if $\Delta^2 u = 0$ in Ω then

$$\|Du\|_{L^\infty(\Omega)} \leq C\|Du\|_{L^\infty(\partial\Omega)}. \quad (3.12)$$

But the constant C is not necessarily 1 and depends on the Lipschitz character of the domain Ω .

Furthermore, to the best of our knowledge it is not known whether an interpolation operator exists, which satisfies all of the axioms (P1), (P2'-P3') and (P4-P8). If it exists it would be a nice generalization of the AMLE model for slope-interpolating operators. In its absence the thin-plate model comes close by satisfying all but the maximum principle (P2'), which is also satisfied by the thin-plate model for certain very regular domains Ω .

3.3 Kriging

A widely used procedure to interpolate data is Kriging, a method originally developed in the framework of Geostatistics. We present here a brief account of this method, and give some details on a specific case to be used in our experiments. General references on the subject are [124], [53], [103]. The surface to interpolate is modeled as a realization of a random field X , of which we know values $X(x_1), \dots, X(x_n)$ at some sites of \mathbb{R}^2 . Through some second order properties of X (usually to be estimated from the data), a “predictor” \hat{X} is constructed, defined as a linear combination of the known values of X :

$$\hat{X}(x) = \sum_{i=1}^n \lambda(x) X(x_i). \quad (3.13)$$

This predictor satisfies two conditions. First, \hat{X} is requested to be unbiased, that is, for all x , writing E for the mathematical expectation,

$$EX(x) = E\hat{X}(x). \quad (3.14)$$

The predictor also minimizes the least square error, that is, for all x ,

$$E \left(\left| \hat{X}(x) - X(x) \right|^2 \right) = \min_{\{\lambda_i\}} E \left(\left| \sum_i \lambda_i X(x_i) - X(x) \right|^2 \right). \quad (3.15)$$

In other words, \hat{X} is a Best Linear Unbiased Predictor (BLUP), a well known predictor in Statistics.

If the field X is second order stationary (EX is constant and $E(X(x)X(y)) = C(x - y)$, for some function C), then \hat{X} is found by inverting a linear system depending only on C and the data $\{x_1, \dots, x_n\}$ (see [53]). However, and this is quite clear in the case of DEM, X cannot be assumed to be stationary. Therefore X is assumed to have a “drift”, that is $E(X(x)) = \sum_j \alpha_j f_j(x)$, for some functions $\{f_j\}$ of \mathbb{R}^2 . The general mathematical framework in which to address the Kriging problem is the one of “intrinsic random functions”, a concept first introduced by Matheron ([125]) following the work of Gel’fand and Vilenkin ([81]). Let us briefly introduce this particular type of random fields. We define P_k to be the set of polynomials of \mathbb{R}^2 of degree less or equal to k , and an increment with respect to P_k to be a set $\{(\beta_j, y_j)\}_{j=1\dots m}$ of $(\mathbb{R} \times \mathbb{R}^2)^m$, for some integer m , such that

$$\sum_{j=1}^m \beta_j f(y_j) = 0,$$

for all functions $f \in P_k$. An intrinsic random function of order k (k -IRF) is a random field such that there exist a function $\gamma : \mathbb{R}^2 \rightarrow \mathbb{R}$, a “generalized covariance”, such that

$$\text{var} \left(\sum_i \beta_i X(y_i) \right) = \sum_{i,j} \beta_i \beta_j \gamma(y_i - y_j),$$

for all increments $\{(\beta_i, y_i)\}$. These fields are completely determined by the generalized covariance γ , up to a polynomial of degree k . Let us write N_k for the dimension of the space P_k , and $\{p_i\}_{i=1 \dots N_k}$ for a basis of this space. In this chapter, we are interested in particular fields of the form:

$$X(x) = \sum_{i=0}^{N_k} a_i p_i(x) + \epsilon(x), \quad (3.16)$$

where ϵ is a k -IRF such that $E\epsilon = 0$, and the a_i 's are random variables. Fields of the type (3.16) may be shown to be k -IRF, and are said to have a polynomial drift. The Kriging problem is then expressed, in this particular case, in the same manner as in the stationary case, that is to say that \hat{X} must satisfy (3.14) and (3.15). It may be shown that this predictor \hat{X} is obtained by inverting a linear system, depending only on γ and the data (x_i) .

We then consider the following function, defined for $\alpha > 0$, which may be shown to be a valid generalized covariance, provided $\text{Int}(\alpha/2) \leq k$, where $\text{Int}(\alpha)$ is the integer part of $\alpha/2$:

$$\gamma(x) = \begin{cases} C_\alpha |x|^\alpha & \text{if } \frac{\alpha}{2} \notin \mathbb{N} \\ C_\alpha |x|^\alpha \log(|x|) & \text{if } \frac{\alpha}{2} \in \mathbb{N} \end{cases}, \quad (3.17)$$

where C_α are convenient constants depending on α . In the remaining of this chapter, we consider Kriging with fields having such a generalized covariance. This is primarily motivated by the fact that, as we will see more precisely, this choice of γ lead to the same solution of the interpolation problem as with the thin-plates spline method, when choosing $\alpha = 2$. We will thus be in a position to compare our results with thin-plates and Kriging interpolators, in a unified manner with various values of α (see Section 4). We further assume that the value of k (the degree of the polynomial drift) is 1, and thus allows for values $0 < \alpha < 4$. In this case, and when $\alpha = 2$, \hat{X} (actually the realization of \hat{X} corresponding to the realization of X) may be shown to be a solution of the biharmonic problem (3.10), on \mathbb{R}^2 , with limit conditions (keeping the same notations as in (3.10)) $u(x_i) = \phi(x_i)$, for $i = 1, \dots, n$, and $\frac{\partial u}{\partial n} = 0$ at infinity (see [126] and [110]). Numerically, the equivalence may be seen by

considering radial basis functions as discussed in Appendix C.3. Explicitly, we write $\mathbf{g} = (\gamma(x_0 - x_1), \dots, \gamma(x_0 - x_n))$, $\mathbf{p} = (p_1(x_0), \dots, p_{N_k}(x_0))$, $\mathbf{z} = (Z(x_1), \dots, Z(x_n))^T$, and define the matrices $\mathbf{G} = ((\gamma(x_i - x_j)))$ and $\mathbf{P} = ((p_j(x_i)))$.⁶ Then, the value of the predictor at a point x_0 is given by (see [53]):

$$\hat{Z}(x_0) = \mathbf{g}\mathbf{a} + \mathbf{p}\mathbf{b}, \quad (3.18)$$

where \mathbf{a} and \mathbf{b} are column vectors such that

$$\begin{aligned} \mathbf{G}\mathbf{a} + \mathbf{P}\mathbf{b} &= \mathbf{z} \\ \mathbf{P}^T\mathbf{a} &= 0. \end{aligned} \quad (3.19)$$

When $k = 1$ and γ is given by (3.17) with $\alpha = 2$, these are the same equations as in the resolution of the thin-plate problem through radial functions as in [37], since in that case the variogram $\gamma(x) = |x|^2 \log(x)$ is the fundamental solution of the biharmonic equation $\Delta^2\gamma = \delta$. More generally, when $\alpha \neq 2$, the Kriging solution may be obtained by using γ as a radial function.

4 Experiments

The following figures and tables show some results of interpolation by both the AMLE model, and the classical thin-plates and Kriging models discussed in this paper. To illustrate these methods we chose the elevation model shown in figure 3.1(a) (where height is represented as a three-dimensional illuminated surface). This is a $12.42 \times 6.9\text{km}$ DEM around Mount Sainte Victoire where each pixel represents a $30 \times 30\text{m}$ patch. Height values vary between 190 and 1011m and have a precision of about 2m. As explained in Section 3.3, we chose a Kriging model with generalized covariance given by Equation (3.17), so that thin-plate interpolation corresponds mathematically to the case $\alpha = 2$. The use of radial functions enables these two methods to agree also from the numerical point of view.

As a first experiment we quantized this elevation model by keeping only the level curves at regular intervals of 50m height, as a way to simulate the kind of data that can be obtained from scanning non-digital cartography, see figure 3.1(b). Then we try to reconstruct the original DEM by interpolation using the AMLE model (see figure 3.2(a)) and the thin-plate models (see figure 3.2(b)).

⁶ Let's recall that in our experiments we only consider 1-IRFs, so $k = 1$, $N_k = 3$, and the basis of P_k is made up of $p_1\left(\frac{x}{y}\right) = 1$, $p_2\left(\frac{x}{y}\right) = x$, and $p_3\left(\frac{x}{y}\right) = y$.

As a second experiment we simulated the kind of interpolation domains that arise in DEMs produced from stereoscopic views. We present four examples: in the first one we eliminate the top of a mountain (figure 3.3(a)), in the following two, we eliminate a closed region on the side of it (figures 3.3(b) and 3.3(c)) and in the last one we eliminate a flat region (figure 3.3(d)). The corresponding interpolations by the Kriging models for different variograms can be observed in figures 3.4, 3.6, 3.8 and 3.10. The corresponding interpolations by the AMLE model can be observed in figures 3.5(b), 3.7(b), 3.9(b) and 3.11(b). In table 3.1 we display the L^2 distances between original and interpolated DEMs for the different methods.

In these experiments we can observe how (unlike the AMLE model), the thin-plate and Kriging models allow to interpolate slopes, and produce C^1 interpolators (*i.e.* with one continuous derivative) across imposed level curves. Thanks to this ability they allow to roughly recreate the mountain top in figure 3.4.

On the other hand, they are somehow much more diffusive than the AMLE model, producing too smooth reliefs. For instance in examples b and c (figures 3.7(b) and 3.9(b)) we see how the AMLE model better preserved the ridges. In the second case AMLE is better both visually and in terms of RMS error, whereas in the first case AMLE is visually better despite a slightly larger RMS error. A similar observation can be made about the last example (d) in figure 3.11(b), where AMLE is visually better despite a slightly larger RMS error. In general, the AMLE produces a better result when the region to be interpolated is entirely contained on one side of the mountain or in a roughly flat region. In those cases the relative slope variations around the boundary are not important so it is not necessary to actually impose them: the average slope is implicit in the boundary data.

Furthermore the fact that thin-plate and Kriging do not satisfy a maximum principle means that they are less safe than the AMLE model, in the sense that it can create unnatural oscillations. With respect to Kriging, we can observe how a larger value of α can be used to reduce the diffusive behaviour of the thin-plate ($\alpha = 2$), and make it better preserve ridges, although not as good as the AMLE. However, larger values of α lead also to increasingly ill-conditioned systems and to even more unnatural oscillations, as it can be observed for instance in figure 3.6(d).

4.1 A simple combination of the AMLE and thin-plate models

These experiments suggest that it would be interesting to combine the slope-interpolating capability of thin-plate and Kriging models with the better ridge preservation of the

AMLE model. Here we show a very simple way to combine them which allows to preserve many of the good properties of both models.

Let's write our interpolator $u = u_1 + u_2$ as a linear combination of two functions u_1 and u_2 which are respectively solutions of the following PDE's:

$$\begin{aligned} D^2 u_1(Du_1, Du_1) &= 0 & \text{for } x \in \Omega \\ u_1(x) &= \varphi(x) & \text{for } x \in \partial\Omega \end{aligned} \quad (3.20)$$

and

$$\begin{aligned} \Delta^2 u_2(x) &= 0 & \text{for } x \in \Omega \\ u_2(x) &= 0, & \text{for } x \in \partial\Omega \\ \frac{\partial u_2}{\partial n}(x) &= \psi(x) - \frac{\partial u_1}{\partial n}(x) & \text{for } x \in \partial\Omega. \end{aligned} \quad (3.21)$$


This means that we interpolate the boundary data φ as usual with AMLE (obtaining u_1). Then we check at the boundary the slope differences $v = \psi - \frac{\partial u_1}{\partial n}$ between the imposed slope ψ and the AMLE slope $\frac{\partial u_1}{\partial n}$. Then we build a correction term by thin-plate interpolation with zero boundary condition and slope condition at the boundary equal to the slope difference v .

Observe that the combined interpolator u actually matches the boundary values $u = \varphi$ and slopes $\frac{\partial u}{\partial n} = \psi$ as the thin-plate, and Kriging. However, it is not necessarily biharmonic, and it inherits to a large extent the ridge-preservation property of AMLE. This effect is quite visible in figure 3.9(c) and even the hill top in figure 3.5(c). Thanks to this combination of the good properties of both methods, the combined method shows also the lowest L^2 interpolation error with respect to the reference in all cases we tested, except in example c, where the AMLE model produced the best result (see table 3.1).


On the down-side, this procedure might create unnatural oscillations so as the thin-plate and Kriging do. This effect can be observed *e.g.* in the interpolation of example b, which is shown in figure 3.7(c). To reduce this effect we might be interested in interpolating with $u = u_1 + \beta u_2$ for some $\beta \in (0, 1)$, thus coming closer to AMLE, at the cost of a worse match of the slope boundary conditions. Alternatively we could use Kriging with a small power like $\alpha = 1$ (instead of $\alpha = 2$ used for the thin-plate), at the cost of a more diffusive behaviour, *i.e.* less ridge preservation.

5 Discussion and Further Work

In this chapter we propose the AMLE as a new method for interpolating digital elevation models and explore its relationship to previously used methods.

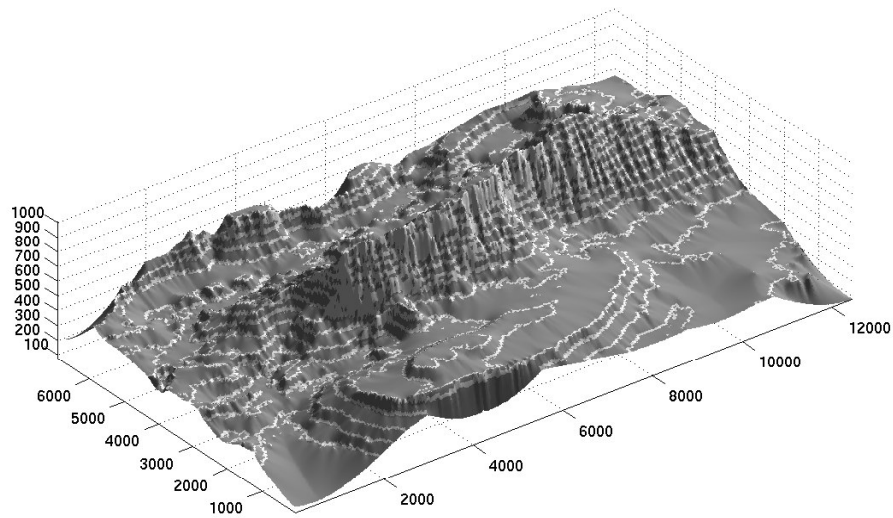


(a) Original DEM

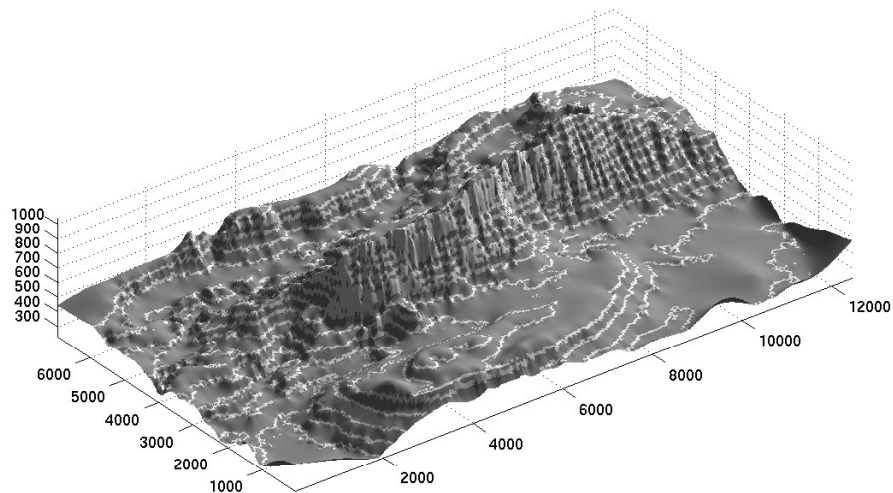


(b) Quantized DEM

Figure 3.1: (a) Elevation model used throughout the experiments. It represents a 12.4×6.9 km area around Mount Sainte Victoire with heights values ranging between 190 and 1011 metres. The horizontal sampling is 30m, whereas the vertical precision is about 2m standard deviation. In all figures all axes values are expressed in metres, but the vertical axis has been stretched by a factor of three with respect to the real aspect ratio in order to better visualize the irregularities of the terrain. (b) In the quantized DEM light gray colors represent known data points whereas dark gray colors represent unknown points to be interpolated. It has been obtained from the original by only keeping level-curves at regular intervals of 50m height.




(a) AMLE interpolation




(b) Thin-plate interpolation


Figure 3.2: Interpolations of the quantized elevation model from figure 3.1b. Light gray colors represent known data points whereas dark gray colors represent interpolated points. Observe how AMLE produces certain flat regions and slope discontinuities, but better preserves ridges and avoids some oscillations produced by the thin-plate spline.




(a) Original DEM for example a



(b) Original DEM for example b




(c) Original DEM for example c




(d) Original DEM for example d


Figure 3.3: Reference data and interpolation domains chosen for the second experiment. The surface represents the ground truth for the elevation model. Light gray colors represent known data points that can be used by the interpolators whereas dark gray colors represent the region to be interpolated.




(a) Reference elevation model in figure 3.3(a)



(b) Kriging interpolation of figure 3.3(a), $\alpha = 1$




(c) Kriging (and thin plates) interpolation of figure 3.3(a), $\alpha = 2$




(d) Kriging interpolation of figure 3.3(a), $\alpha = 3$


Figure 3.4: Interpolations of the first example, figure 3.3(a) (to be continued in the next figure). *Legend:* Light gray colors represent known data points whereas dark gray colors represent interpolated points.



(a) Reference elevation model in figure 3.3(a)




(b) AMLE interpolation of figure 3.3(a)




(c) AMLE + thin-plate interpolation of figure 3.3(a)


Figure 3.5: Interpolations of the first example, figure 3.3(a) (continued from previous figure). Observe how the bad performance of AMLE is corrected by the combined method, while still preserving ridges better than the Kriging methods. *Legend:* Light gray colors represent known data points whereas dark gray colors represent interpolated points.




(a) Reference elevation model in figure 3.3(b)



(b) Kriging interpolation of figure 3.3(b), $\alpha = 1$




(c) Kriging (and thin plates) interpolation of figure 3.3(b), $\alpha = 2$




(d) Kriging interpolation of figure 3.3(b), $\alpha = 3$


Figure 3.6: Interpolations of the second example, figure 3.3(b) (to be continued in the next figure). Among Kriging methods, ridges are best preserved by the highest values of α , but at the cost of more oscillations than lower values of α and AMLE. Lower values of α , on the other hand are much too diffusive. *Legend:* Light gray colors represent known data points whereas dark gray colors represent interpolated points.



(a) Reference elevation model in figure 3.3(b)




(b) AMLE interpolation of figure 3.3(b)




(c) AMLE + thin-plate interpolation of figure 3.3(b)


Figure 3.7: Interpolations of the second example, figure 3.3(b) (continued from the previous figure). Observe how AMLE (and to some extent the combined AMLE + thin-plate method) better preserve the ridges than Kriging methods. The combined method, however, may introduce unnatural oscillations, so as Kriging and thin-plate methods do. *Legend:* Light gray colors represent known data points whereas dark gray colors represent interpolated points.




(a) Reference elevation model in figure 3.3(c)



(b) Kriging interpolation of figure 3.3(c), $\alpha = 1$




(c) Kriging (and thin plates) interpolation of figure 3.3(c), $\alpha = 2$




(d) Kriging interpolation of figure 3.3(c), $\alpha = 3$


Figure 3.8: Interpolations of the third example, figure 3.3(c) (to be continued in the next figure). *Legend:* Light gray colors represent known data points whereas dark gray colors represent interpolated points.



(a) Reference elevation model in figure 3.3(c)



(b) AMLE interpolation of figure 3.3(c)



(c) AMLE + thin-plate interpolation of figure 3.3(c)

Figure 3.9: Interpolations of the third example, figure 3.3(c). (continued from the previous figure). Observe how AMLE preserves the ridges remarkably better than Kriging. *Legend:* Light gray colors represent known data points whereas dark gray colors represent interpolated points.

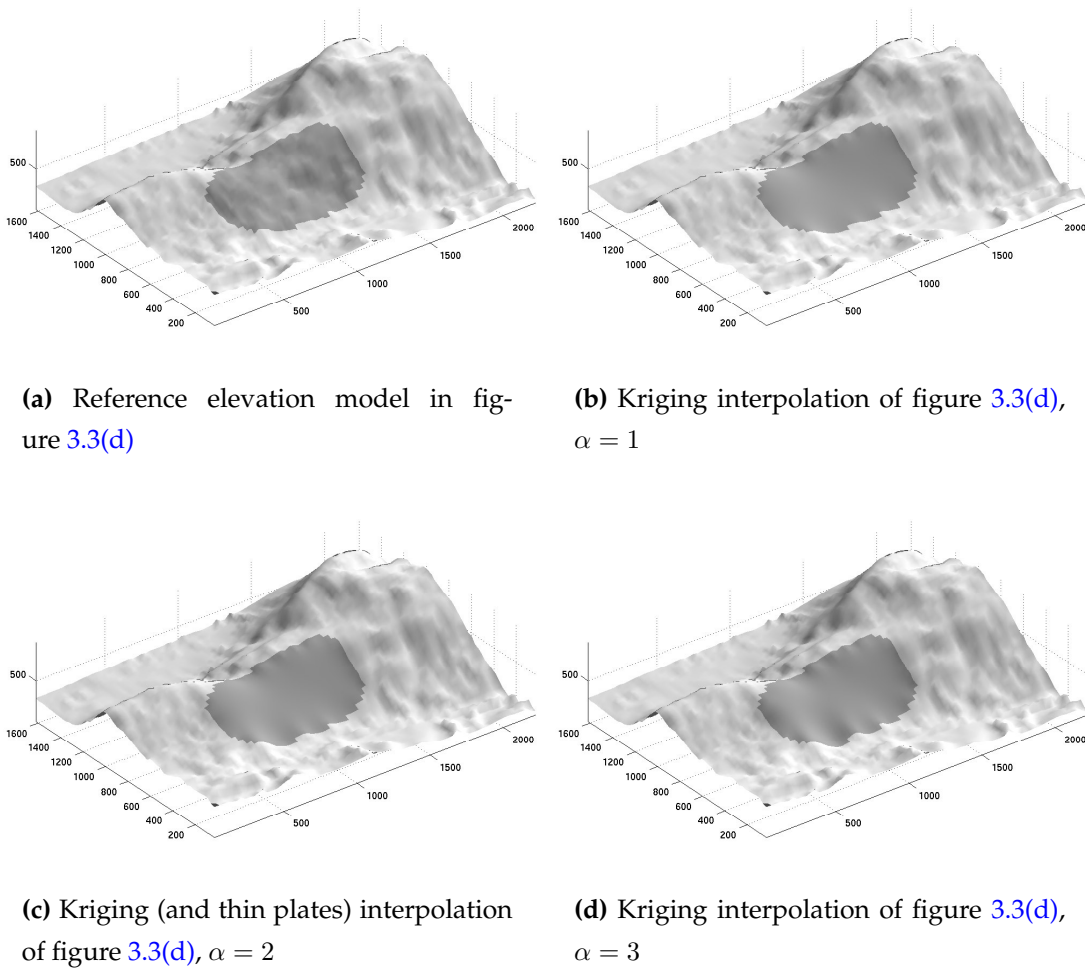




Figure 3.10: Interpolations of the fourth example, figure 3.3(d) (to be continued in the next figure). *Legend:* Light gray colors represent known data points whereas dark gray colors represent interpolated points.

method	example a	example b	example c	example d
AMLE	52.88	25.21	25.03	10.64
Kriging, $\alpha = 1$	26.57	28.22	37.22	7.15
Kriging, $\alpha = 2$ (thin plates)	16.72	24.84	37.00	5.41
Kriging, $\alpha = 3$	20.57	22.40	32.41	5.81
Kriging, $\alpha = 3.9$	27.94	25.90	29.01	8.75
AMLE + thin plate	13.98	21.87	35.82	5.37


Table 3.1: L^2 distances between originals and interpolated for examples a-d, figure 3.3 (all values are expressed in metres)



(a) Reference elevation model in figure 3.3(d)



(b) AMLE interpolation of figure 3.3(d)



(c) AMLE + thin-plate interpolation of figure 3.3(d)

Figure 3.11: Interpolations of the fourth example, figure 3.3(d). (continued from the previous figure). *Legend:* Light gray colors represent known data points whereas dark gray colors represent interpolated points.

method	example a	example b	example c	example d
AMLE	0.2 (1.97)	1.1 (15.36)	3.9 (20.49)	1.1 (11.77)
Kriging, $\alpha = 1$	0.14 (2.78)	1.84 (22.78)	1.74 (20.58)	0.69 (9.45)
Kriging, $\alpha = 2$ (thin plates)	0.15 (2.89)	2.23 (23.47)	1.97 (21.72)	0.86 (9.78)
Kriging, $\alpha = 3$	0.16 (2.78)	2.11 (22.78)	1.74 (20.58)	0.87 (9.45)
Kriging, $\alpha = 3.9$	0.17 (2.78)	1.94 (22.78)	1.20 (20.58)	0.82 (9.45)
N	197	768	745	471
M	140	296	285	216

Table 3.2: Computational resources for the examples in table 3.1. All values are expressed in seconds (and Mflops). At the end we indicate the size of the problem in terms of N (size of region to be interpolated in pixels) and M (size of region boundary, *i.e.* the number of interpolation constraints).

Unlike other methods that have been proposed in the geoscience literature, the AMLE satisfies a maximum principle, which ensures that it does not create oscillations. We also showed through experiments that it also preserves ridges much better than Kriging and thin-plate methods, which are much more diffusive. Even though Kriging for high values of α comes a bit closer to the AMLE (in terms of ridge preservation), it's exactly when α is high that Kriging produces the most undesirable oscillations. Also in terms of computational efficiency it can be faster than Kriging, depending on the data.

In the case of interpolation of iso-level lines AMLE can be interpreted (under certain regularity hypothesis) as iterated weighted geodesic distance interpolators, which explains why it shares some qualitative properties with this method. Furthermore, since this iteration converges faster than numerical schemes for solving AMLE, it can be used to accelerate these schemes.

The major disadvantage of AMLE is the fact that it cannot interpolate slopes as well as values. For this purpose we proposed a reasonable generalization to slope-interpolating operators of the axiomatic approach that leads to the AMLE model. None of the currently known methods satisfies all of these axioms. The thin-plate method comes closest by satisfying all but the maximum principle, hence it may produce unnatural oscillations. The search for a (possibly non-linear) fourth-order differential operator satisfying all of the proposed axioms remains an interesting problem for future research, since it might lead to a slope-interpolating operator with all the advantages of AMLE.

The discussion in this chapter suggests to try fourth order differential operators satisfying a differential equation like

$$D^4u(q(u), q(u), q(u), q(u))(x) = 0, \quad \text{for } x \in \Omega \quad (3.22)$$

or minimizing the variational principle

$$\int_{\Omega} |D^2u(q(u), q(u))(x)|^p dx \quad (3.23)$$

both subject to the usual boundary conditions

$$u(x) = \varphi(x), \quad \text{for } x \in \partial\Omega \quad (3.24)$$

$$\frac{\partial u}{\partial n}(x) = \psi(x) \quad \text{for } x \in \partial\Omega. \quad (3.25)$$

In both cases $q(u)$ represents the eigenvector of the Hessian matrix D^2u which corresponds to the eigenvalue closest to 0

$$q(u) = \arg \min_{q \in \text{Ker}(D^2u - \lambda I), \|q\|=1} |D^2u(q, q)|, \quad (3.26)$$

i.e. the orientation of minimal curvature of u . Put another way $q(u)$ represents the orientation of the ridges, and the PDE or variational principle just described, mean that we prefer to diffuse in the direction of the ridges, rather than across ridges.

However, such fourth order differential equations are usually very unstable numerically and it is difficult to obtain existence and uniqueness results under reasonable hypothesis.

An alternative could be an adaptation of the model proposed in [26] for image disocclusion. In this model the image gray-levels u and the gradient orientations $\theta = \frac{Du}{|Du|}$ are jointly propagated inside Ω in such a way that they match the boundary conditions, while the interpolated vector flow θ is regular and the interpolated image u follows this vector flow, *i.e.* $\theta \approx \frac{Du}{|Du|}$. This procedure is good for image disocclusion because it allows to smoothly propagate edges inside the occlusion.

When interpolating elevation models, the main characteristics we want to preserve are ridges rather than edges. Therefore an adaptation of this disocclusion procedure, where we jointly propagate the elevation u and the ridge orientation $\theta = q(u)$, while the interpolated vector flow θ is regular and the interpolated elevation u follows this vector flow, *i.e.* $\theta \approx q(u)$, would be of interest for the interpolation of elevation data.

In the absence of such a generalized interpolator, which satisfies all desired axioms, and in addition shows better ridge preservation, we can use a combination of AMLE

and Kriging or thin-plate methods. We showed here how a simple combination of such methods did in many cases better than all known methods, both in terms of L^2 differences and in terms of visual quality. The good results obtained by such a simple combination suggests that there is still much room for improvement by a smarter combination of known methods. Eventually, manual tuning of the α and β parameters for the proposed combined method at each connected component of the interpolation domain, may be required to obtain the best results.

Appendix

A AMLE and weighted geodesic distance interpolations

A.1 Review of weighted geodesic distance interpolation

For the sake of completeness we recall here the precise definition of weighted geodesic distances as defined in [164]. They are defined in three steps:

Definition 9 (Geodesic time, curves and distance). *The weighted geodesic time to reach a set $A \subseteq \Omega$ from a point $x \in \Omega$ within a region $\Omega \subseteq \mathbb{R}^2$ with potential $g : \Omega \rightarrow \mathbb{R}^+$ is defined as*

$$t_{\Omega,g}(x, A) = \min_{\gamma} \int_0^1 g(\gamma(p)) |\gamma'(p)| dp \quad (3.27)$$

where the minimization is performed over all continuous curves $\gamma : [0, 1] \rightarrow \Omega$ joining x and A , i.e. such that $\gamma(0) = x$ and $\gamma(1) \in A$.

The corresponding set $G_{\Omega,g}(x, A)$ of geodesic curves or simply geodesics is the set of minimizers of equation (3.27):

$$G_{\Omega,g}(x, A) = \arg \min_{\gamma} \int_0^1 g(\gamma(p)) |\gamma'(p)| dp. \quad (3.28)$$

Finally the geodesic distance is the minimal geometric length of a geodesic curve:

$$d_{\Omega,g}(x, A) = \min_{\gamma \in G_{\Omega,g}(x, A)} \int_0^1 |\gamma'(p)| dp. \quad (3.29)$$

The members of the set $\hat{G}_{\Omega,g}(x, A)$ of minimizers of (3.29) are called minimal length geodesics.

For convenience we shall also use the notation $\tilde{\gamma}$ for the arc-length parameterization of the curve γ , such that $\tilde{\gamma}(s) = \gamma(p)$ with $ds = |\gamma'(p)| dp$, so that $|\tilde{\gamma}'(s)| = 1$ and $\langle \tilde{\gamma}'', \tilde{\gamma}' \rangle = 0$, with $|\tilde{\gamma}''(s)| = \kappa$ equal to the curvature of γ .

Based of the geodesic distance we can define a geodesic distance interpolator with potential g as in equation (3.6). Then a fixed point is a geodesic distance interpolator whose gradient module is the same potential we used in the interpolation.

Definition 10 (Geodesic distance interpolator and fixed point). *Given a region Ω whose boundary can be partitioned in two disjoint regions $\partial\Omega = \partial\Omega_1 \cup \partial\Omega_2$, and the boundary conditions $u|_{\partial\Omega_1} = c_1 \in \mathbb{R}$ and $u|_{\partial\Omega_2} = c_2 \in \mathbb{R}$, a geodesic distance interpolator with potential g is defined as*

$$u(x) = E_{\Omega,g}(c_1, c_2)(x) = \frac{c_2 L_1(x) + L_2(x) c_1}{L_1(x) + L_2(x)} \quad \text{where} \quad L_i(x) = d_{\Omega,g}(x, \partial\Omega_i). \quad (3.30)$$

Such an interpolator is called a fixed point if

$$u = E_{\Omega,g}(c_1, c_2) \quad \text{with} \quad g(x) = |Du(x)| \quad (3.31)$$

for all $x \in \Omega$.

A.2 Equivalence with AMLE in the twice differentiable case

Our first two results are not the best way to show the equivalence between fixed points and AMLE, but they provided a convenient way to deduce the weight needed in equation (3.30) to have an AMLE.

Proposition 6. *If $\gamma \in G_{\Omega,g}(x, A)$ is a twice differentiable geodesic and the potential g is also differentiable, then we have*

$$\langle Dg, \tilde{\gamma}'' \rangle - g|\tilde{\gamma}''|^2 = 0. \quad (3.32)$$

If in addition $g(x) \neq 0$ for $x \in A$, any geodesic arrives normal to the set A , i.e. for some $k \in \mathbb{R}$

$$\gamma'(1) = kn_A \quad (3.33)$$

where n_A is normal to the set A .

Proof. We consider any admissible curve $\gamma : [0, 1] \rightarrow \Omega$ joining x and A , i.e. such that $\gamma(0) = x$ and $\gamma(1) \in A$. A variation $\gamma + \delta\gamma$ to this curve is still admissible if $\delta\gamma(0) = 0$ and $\delta\gamma(1)$ is tangent to A . The variation of the geodesic time becomes:

$$\begin{aligned} \delta t &= -t + \int_0^1 g(\gamma(p) + \delta\gamma(p)) |\gamma'(p) + \delta\gamma'(p)| dp \\ &= \int_0^1 \langle \delta\gamma(p), Dg(\gamma(p)) \rangle |\gamma'(p)| + \frac{\langle \delta\gamma', \gamma' \rangle}{|\gamma'|} g(\gamma(p)) dp + o(\delta\gamma) \end{aligned}$$

after a Taylor development of u and $x \mapsto |x|$ around $\gamma(p)$. Integrating by parts and reparameterizing γ with the arc-length parameter we get

$$\delta t = [g(\gamma(p)) \langle \tilde{\gamma}'(s(p)), \delta\gamma(p) \rangle]_{p=0}^{p=1} + \int_0^L \left\langle \delta\gamma, \left[Dg(\tilde{\gamma}(s)) - \frac{d}{ds} \left(g(\tilde{\gamma}(s)) \tilde{\gamma}'(s) \right) \right] \right\rangle ds + o(\delta\gamma).$$

To minimize the geodesic time we should have $\delta t = o(\delta\gamma)$ for any admissible small variation $o(\delta\gamma)$. So, from the previous equation, we deduce that for γ to be a geodesic we should have: first $\gamma'(1)$ orthogonal to $\delta\gamma(1)$ and hence tangent to A , whenever $g(\gamma(1)) \neq 0$, and second

$$\begin{aligned} Dg &= \frac{d}{ds} \left(g(\tilde{\gamma}(s)) \tilde{\gamma}'(s) \right) \\ &= (\langle \tilde{\gamma}', Dg \rangle) \tilde{\gamma}' - g \tilde{\gamma}'' . \end{aligned}$$

Multiplying the last equation by $\tilde{\gamma}'$ yields $\langle \tilde{\gamma}', Dg \rangle = \langle \tilde{\gamma}', Dg \rangle$, but multiplying by $\tilde{\gamma}''$ yields

$$\langle Dg, \tilde{\gamma}' \rangle - g |\tilde{\gamma}''|^2 = 0.$$

□

In the next proposition we need to assume that the pair of (upwards and downwards) geodesics from a point do not change when we move along the geodesic. In that case, geodesic distance interpolators are affine along geodesics, which allows to obtain a differential equation linking the direction of geodesics and the interpolator. Later we shall state conditions under which this hypothesis of geodesics stability is satisfied.

Definition 11 (Globally stable geodesics). *A pair (γ_1, γ_2) of minimal length geodesics to level curves $\gamma_i \in \hat{G}_{\Omega, g}(x, \chi_{u=c_i})$ where $c_1 \leq u(x) \leq c_2$ is called globally stable, if for any $y \in \text{Range}(\gamma_1) \cup \text{Range}(\gamma_2)$, there is a minimal length geodesic pair (δ_1, δ_2) from y to the same level curves $\delta_i \in \hat{G}_{\Omega, g}(y, \chi_{u=c_i})$, such that*

$$\text{Range}(\delta_1) \cup \text{Range}(\delta_2) = \text{Range}(\gamma_1) \cup \text{Range}(\gamma_2). \quad (3.34)$$

Put another way, the geodesics are globally stable, if they do not change when we move along the geodesic.

Proposition 7. *If $u = E_{\Omega, g}(c_1, c_2)$ is a twice differentiable weighted geodesic distance interpolator, such that the pairs (γ_1, γ_2) of geodesics $\gamma_i \in G_{\Omega, g}(x, \chi_{u=c_i})$ used in the interpolation are globally stable, then we have*

$$\langle Du, \tilde{\gamma}_i'' \rangle + D^2u(\tilde{\gamma}_i', \tilde{\gamma}_i') = 0. \quad (3.35)$$

Proof. Observe that under the hypothesis of geodesics stability $L_1(x) + L_2(x) = L$ measures the total length of the stable geodesic through x , which has a constant length L when x moves along the geodesic.

Now consider a point x_h at distance h from x measured along the geodesic $\gamma_1 \in G_{\Omega, g}(x, \chi_{u=c_1})$, *i.e.*

$$x_h = \tilde{\gamma}_1(h).$$

Then $L_1(x_h) = L_1(x) - h$ and $L_2(x_h) = L_2(x) + h$. Thus

$$\frac{u(x_h) - u(x)}{h} = \frac{c_1 - c_2}{L}$$

and we deduce that $\frac{du(\tilde{\gamma}_1(s))}{ds} = \frac{c_1 - c_2}{L}$ at $s = 0$. Similarly we obtain $\frac{du(\tilde{\gamma}_2(s))}{ds} = \frac{c_2 - c_1}{L}$ at $s = 0$ (the sign changes because the γ_1 and γ_2 go in opposite directions).

Now consider the upwards concatenation of γ_1 and γ_2 , *i.e.* $\gamma : [0, 1] \rightarrow \Omega$ such that $\gamma(0) = \gamma_1(0)$, $\gamma(1) = \gamma_2(1)$, and $\text{Range}(\gamma) = \text{Range}(\gamma_1) \cup \text{Range}(\gamma_2)$. Since the geodesic pairs are globally stable, the previous argument actually shows that $\frac{d}{ds}u(\tilde{\gamma}(s)) = \frac{c_2 - c_1}{L}$ all along the curve, and that this value is constant. Hence $\frac{d^2}{ds^2}u(\tilde{\gamma}(s)) = 0$.

The rest of the proof consists of simply applying the chain rule to this last expression, to obtain equation (3.35). \square

Proposition 7 suggests that a simple way for a geodesic distance interpolator to satisfy AMLE is that $\tilde{\gamma}' = \frac{Du}{|Du|}$, so that the first term in equation (3.35) becomes zero and the equation reduces to $\frac{1}{|Du|^2}D^2u(Du, Du) = 0$.

Now we wonder how to deduce g such that $\tilde{\gamma}' = \frac{Du}{|Du|}$. To do so, we first observe that this property, along with proposition 7 implies that $\tilde{\gamma}'' = \frac{\langle D^2u, Du \rangle}{|Du|^2}$, and we substitute in proposition 6 (to simplify the computations we use just a sufficient condition for proposition 1) and we obtain

$$Dg = g\tilde{\gamma}'' = g \frac{\langle D^2u, Du \rangle}{|Du|^2}. \quad (3.36)$$

A simple way to satisfy this equation (and hence proposition 1) is to set $g = |Du|$.

This argument shows that for this choice of potential $g = |Du|$, if the geodesics are globally stable and if they follow the direction of Du , then the corresponding geodesic distance interpolator, if twice differentiable is AMLE. But this argument doesn't show that for this choice of potential we actually have $\tilde{\gamma}' = \frac{Du}{|Du|}$. Note that proposition 6 does not uniquely determine geodesics from the potential g , it only imposes a 1D local constraint on them. The actual geodesics are chosen globally and also depend on the set A .

The surprising fact, to be established in the next proposition, is that for our choice of potential $g = |Du|$, and boundary conditions (A is a level curve of u), the geodesics are actually *locally* determined if u is C^1 , namely by the gradient direction.

Proposition 8. *Let $u : \Omega \rightarrow \mathbb{R}$ be differentiable inside Ω , and the potential be $g = |Du|$. Then the geodesic time to go from one level curve to another level-curve of u is independent of the position of the starting point on the first level curve. More precisely, let $x \in \Omega$, and $\chi_{u=c}$ be any non-empty level-curve of u . Then*

$$t_{\Omega,g}(x, \chi_{u=c}) = |c - u(x)|. \quad (3.37)$$

Furthermore, the geodesic $\gamma \in G_{\Omega,g}(x, \chi_{u=c})$ attaining this geodesic time satisfies

$$\tilde{\gamma}' = \sigma \frac{Du}{|Du|} \quad \text{where} \quad \sigma = \text{sign}(c - u(x)) \quad (3.38)$$

at almost every point where $|Du| \neq 0$. Therefore, if u is C^1 , equation (3.38) holds at every point where $|Du| \neq 0$.

Proof. The first part follows directly from Cauchy-Schwartz inequality

$$|D(u(\gamma(p)))| |\gamma'(p)| \geq \langle \sigma Du(\gamma(p)), \gamma'(p) \rangle = \sigma \frac{d}{dp} u(\gamma(p)).$$

which becomes an equality only if $|Du| = 0$ or $\gamma' = k\sigma Du$ for some $k > 0$. Equivalently, it becomes an equality iff $|Du| = 0$ or $\tilde{\gamma}' = \sigma \frac{Du}{|Du|}$. Now we integrate on both sides with respect to p to obtain

$$\int_0^1 |D(u(\gamma(p)))| |\gamma'(p)| dp \geq \sigma(c - u(x)) = |c - u(x)|$$

And this lower bound is attained iff $\tilde{\gamma}' = \sigma \frac{Du}{|Du|}$ at almost every point where $Du \neq 0$. Thus the geodesic time attains the lower bound $|c - u(x)|$. \square

The second part of this result can also be shown in another way, which better explains the dynamic programming principle behind this proposition. This is important for the implementation of this scheme. Indeed the first part allows us to write:

$$t_{\Omega,g}(x, \chi_{u=c}) = t_{\Omega,g}(x, \chi_{u=c'}) + t_{\Omega,g}(x', \chi_{u=c}) \quad (3.39)$$

for any intermediate level c' such that $\sigma u(x) < \sigma c' < \sigma c$, and for any intermediate point $x' \in \chi_{u=c'}$ in the intermediate level curve. This means that the geodesic time minimization can be performed independently between each pair of level curves. Put another

way, the set of geodesics at each point can be constructed by dynamic programming. Now consider a geodesic $\gamma \in G_{\Omega,g}(x, \chi_{u=c})$ and an intermediate point $x' = \gamma(p)$ on the geodesic. Then equation (3.39) shows that $\gamma_1 = \gamma|_{[0,p]} \in G_{\Omega,g}(x, \chi_{u=c'})$ and $\gamma_2 = \gamma|_{[p,1]} \in G_{\Omega,g}(x', \chi_{u=c})$. Now applying the second part of proposition 6 to γ_1 we deduce that at the point $x' = \gamma(p)$, the geodesic γ is orthogonal to the level curve $\chi_{u=c'}$, or equivalently, that it is parallel to the gradient of u , i.e. $\tilde{\gamma}'(p) = \sigma \frac{Du}{|Du|}(\gamma(p))$

The main result is a direct consequence of propositions 7 and 8:

Corollary 1. *Assume that the interpolation domain Ω and “level-curve” boundary conditions are chosen as in definition 10, with $c_1 < c_2$.*

If u is a twice differentiable fixed point, then it is AMLE. Conversely if u is a twice differentiable AMLE it is a fixed point.

Proof. If u is twice differentiable (either AMLE or fixed point), it is in particular C^1 , and the geodesics to level curves of u that are obtained with potential $g = |Du|$, satisfy, according to proposition 8, either $\tilde{\gamma}' = \sigma \frac{Du}{|Du|}$ or $Du(\gamma(p)) = 0$ everywhere along the geodesic.

The case $Du(\gamma(p)) = 0$ is excluded either from the fact that u is a geodesic distance interpolator (with $c_1 < c_2$), or a twice differentiable AMLE with $\text{Lip}(u, \partial\Omega) > 0$.

In either case since u is twice differentiable, $\tilde{\gamma}' = \sigma \frac{Du}{|Du|}$, γ is also twice differentiable and we can apply the chain rule to obtain:

$$\frac{d^2}{ds^2}u(\tilde{\gamma}(s)) = \langle Du, \tilde{\gamma}_i'' \rangle + D^2u(\tilde{\gamma}_i', \tilde{\gamma}_i').$$

Now if u is a fixed point, $\tilde{\gamma}' = \sigma \frac{Du}{|Du|}$ everywhere, means that geodesics are globally stable, so using proposition 7 we get $\frac{d^2}{ds^2}u(\tilde{\gamma}(s)) = 0$, and the first term is zero (because Du is parallel to $\tilde{\gamma}'$). Hence $D^2u(\tilde{\gamma}_i', \tilde{\gamma}_i') = \frac{1}{|Du|^2}D^2u(Du, Du) = 0$, so, u is AMLE.

Conversely if u is AMLE then $D^2u(\tilde{\gamma}_i', \tilde{\gamma}_i') = \frac{1}{|Du|^2}D^2u(Du, Du) = 0$, so $\frac{d^2}{ds^2}u(\tilde{\gamma}(s)) = 0$ (because the first term is zero since Du is parallel to $\tilde{\gamma}'$). This means that u is affine along the globally stable geodesics, so u is fixed point. \square

A.3 Partial results in the weakly differentiable case

In the previous section we showed the equivalence between fixed points and AMLE in the case where both are twice differentiable. However this hypothesis is not well justified, since the AMLE is not necessarily regular. The most widespread conjecture is that the AMLE is $C^{1,1/3}$ (C^1 , with Hölder regularity of the first derivative of order $\frac{1}{3}$,

i.e. $|Du(x) - Du(y)| \leq |x - y|^{1/3}$, but even this result has not been proved. An example for this behaviour is the function $u(x, y) = |x|^{4/3} - |y|^{4/3}$, which is $C^{1,1/3}$ but not C^2 .

In this section we try to generalize the results in the previous section to the weakest possible case. In particular we show that a fixed point which is C^1 except at a non-dense set of isolated points is AMLE. However, the converse seems not to be true. Indeed our experiments suggest that at points where the AMLE is not C^2 , the fixed points tend to become non-differentiable (in the sense that $|Du|$ is still defined and non-null, since Lipschitz, but the orientation of the gradient isn't). The level curves of the fixed points tend to show angles at such regions.

We start by generalizing the concept of gradient and gradient module for a Lipschitz image u which is not necessarily differentiable everywhere.

Definition 12 (maximal and minimal slope). *The maximal slope of $u : \Omega \rightarrow \mathbb{R}$ at a point $x \in \Omega$ is*

$$|D^+u|(x) = \limsup_{r \rightarrow 0} \sup_{|y|=r} \frac{u(x+y) - u(x)}{r} \quad (3.40)$$

and the minimal slope is

$$|D^-u|(x) = - \liminf_{r \rightarrow 0} \inf_{|y|=r} \frac{u(x+y) - u(x)}{r} \quad (3.41)$$

The upwards gradient of u is a set of vectors d with module equal to the maximal slope and with a direction along which u locally reaches the maximal slope:

$$D^+u(x) = \left\{ |D^+u|(x) d : |d| = 1 \wedge |D^+u|(x) = \lim_{r \rightarrow 0^+} \frac{u(x+dr) - u(x)}{r} \right\} \quad (3.42)$$

and similarly we define the downwards gradient

$$D^-u(x) = \left\{ |D^-u|(x) d : |d| = 1 \wedge |D^-u|(x) = - \lim_{r \rightarrow 0^+} \frac{u(x+dr) - u(x)}{r} \right\}. \quad (3.43)$$

As an example imagine that u is a pyramid of height 1 and square basis, where the square has side 2. Assume that we place the coordinate system parallel to the sides of the square, centered at the center of the basis and that x belongs to one of the edges of the pyramid (the one placed in the first quadrant for instance). In that case we get:

$$\begin{aligned} |D^+u| &= \sqrt{2}/2 = 1/\sqrt{2} \\ |D^-u| &= 1 \\ D^+u &= \left\{ \frac{1}{2}(-1, -1) \right\} \\ D^-u &= \left\{ (0, 1), (1, 0) \right\} \end{aligned}$$

In addition we say that the gradient module is discontinuous at x if $|D^+u|(x) \neq |D^-u|(x)$. Note that in the previous example the gradient module is discontinuous at the edges. In the classical sense, however, the gradient module would be undefined at the edges and would be equal to 1 everywhere else. Thus in the classical sense the gradient module would admit a continuous completion, which is not the case in the sense we just defined.

Now we can redefine the geodesic time even in the case where the gradient is not defined on a set of zero measure in the 1D sense (like the edge of the pyramid in the previous example). However we have to restrain ourselves to the particular potential $g = |Du|$ and level-curve-like sets A which are entirely above or below the starting point x . The prototype for such sets will be the level curves $A = \chi_{u=c}$.

Definition 13 (Geodesic time, curves and distance). *The weighted geodesic time to reach a set $A \subseteq \Omega$ from a point $x \in \Omega$ within a region $\Omega \subseteq \mathbb{R}^2$ with potential $g : \Omega \rightarrow \mathbb{R}^+$ is defined as*

$$t(x, A) = \begin{cases} \min_{\gamma} \int_0^1 |D^+u|(\gamma(p)) |\gamma'(p)| dp & \text{if } u(A) \geq u(x) \\ \min_{\gamma} \int_0^1 |D^-u|(\gamma(p)) |\gamma'(p)| dp & \text{if } u(A) \leq u(x) \\ \text{undefined} & \text{if } \exists a, b \in A, u(a) < u(x) < u(b) \end{cases} \quad (3.44)$$

where the minimization is performed over all continuous curves $\gamma : [0, 1] \rightarrow \Omega$ joining x and A , i.e. such that $\gamma(0) = x$ and $\gamma(1) \in A$.⁷

The corresponding set $G(x, A)$ of geodesic curves or simply geodesics is the set of minimizers of equation (3.44).

Finally the geodesic distance is the geometric length of the shortest geodesic curve:

$$d(x, A) = \min_{\gamma \in G(x, A)} \int_0^1 |\gamma'(p)| dp. \quad (3.45)$$

The members of the set $\hat{G}(x, A)$ of minimizers of (3.45) are called minimal length geodesics.

With this definition we can generalize proposition 8 on the lower bound for the geodesic time and the orientation of geodesics. We need first a technical lemma.

Lemma 8. *Consider a point $x \in \Omega$ and a level $c > u(x)$. Consider a geodesic $\gamma \in G(x, \chi_{u=c})$ from x to the level curve.*

⁷ Note that this definition is consistent with the previous one, since they coincide whenever the gradient module is continuous, i.e. when $|D^+u| = |D^-u| = |Du|$.

Then for any member $d^+ \in D^+u(\gamma(p))$ of the upwards gradient

$$|D^+u|(\gamma(p)) |\gamma'(p)| \geq \langle d^+, \gamma'(p) \rangle \quad (3.46)$$

with equality iff $|D^+u|(\gamma(p)) = 0$ or $\frac{\gamma'(p)}{|\gamma'(p)|} = \frac{d^+}{|D^+u|(\gamma(p))}$.

If in addition, $\frac{\gamma'(p)}{|\gamma'(p)|} = \frac{d^+}{|D^+u|(\gamma(p))}$, and u is Lipschitz, then

$$\langle d^+, \gamma'(p) \rangle = \frac{d}{dp} u(\gamma(p)). \quad (3.47)$$

Similarly, the same result holds for a lower level $c < u(x)$, by substituting D^+ and d^+ by D^- and d^- .

Proof. The first part of the lemma is just a simple application of Cauchy-Schwartz.

For the second part consider $d_r = \gamma(p+r) - \gamma(p)$, so that

$$\frac{u(\gamma(p+r)) - u(\gamma(p))}{r} = \frac{u(\gamma(p) + d_r) - u(\gamma(p))}{|d_r|} \left(\left\langle \frac{d_r}{|d_r|}, \frac{\gamma(p+r) - \gamma(p)}{r} \right\rangle \right)$$

Taking the limit when $r \rightarrow 0$ the left term tends to $\frac{d}{dp} u(\gamma(p))$ and the scalar product in the right term tends to $\left\langle \frac{d^+}{|d^+|}, \gamma'(p) \right\rangle$. To establish the second part of the theorem we only need to show that

$$\lim_{r \rightarrow 0} \frac{u(\gamma(p) + d_r) - u(\gamma(p))}{|d_r|} = |D^+u|(\gamma(p)).$$

We already know from the definition of upwards gradient that

$$\lim_{r \rightarrow 0} \frac{u(\gamma(p) + d^+r) - u(\gamma(p))}{|d^+|r} = |D^+u|(\gamma(p)).$$

In addition, if C is the Lipschitz constant of u , using a Taylor development of γ we obtain:

$$\begin{aligned} u(\gamma(p) + d_r) &= u(\gamma(p) + r\gamma'(p) + o(r)) \\ &= u(\gamma(p) + r\gamma'(p)) + C o(r) \\ &= u(\gamma(p) + r'd^+) + o(r') \end{aligned}$$

where in the last line we used the hypothesis $\frac{\gamma'(p)}{|\gamma'(p)|} = \frac{d^+}{|D^+u|(\gamma(p))}$, and $r' = |\gamma'(p)|r$. Finally

$$\begin{aligned} \lim_{r \rightarrow 0} \frac{u(\gamma(p) + d_r) - u(\gamma(p))}{|d_r|} &= \lim_{r \rightarrow 0} \frac{u(\gamma(p) + rd^+) - u(\gamma(p))}{r|d^+|} \\ &= |D^+u|(\gamma(p)). \end{aligned}$$

□

Proposition 9. *If u is Lipschitz then the geodesic time from x to a level curve at level c is*

$$t(x, \chi_{u=c}) = |c - u(x)|. \quad (3.48)$$

If in addition the level $c > u(x)$ is higher than $u(x)$, then any upwards geodesic $\gamma \in G(x, \chi_{u=c})$ attaining this lower bound satisfies

$$\gamma'(p) = kd^+ \quad \text{for some } k > 0, \text{ and } d^+ \in D^+u(\gamma(p)) \quad (3.49)$$

for almost every p where $\gamma(p)$ is differentiable, and $|D^+u|(\gamma(p)) \neq 0$.

Similarly for downwards geodesics, substituting D^+ and d^+ by D^- and d^- .

Proof. After applying the lemma, the proof follows exactly as in proposition 8 □

Finally we generalize proposition 7 which establishes that fixed points are affine along the geodesics with slope $\frac{c_2 - c_1}{L}$. Here, we assume weaker hypothesis, and the fixed points are no longer affine, but close enough to affine to allow us to show that they are also AMLE.

When accepting less regular u we have to also accept that at certain points D^+u contains more than one direction as in the case of the edge of the pyramid presented at the beginning of this section. This means that at such points the minimization of the geodesic distance might choose one direction or the other depending on the global length.

Thus geodesics can no longer be assumed globally stable, but a weaker stability assumption turns out to be sufficient.

Definition 14 (Locally stable geodesics). *A pair (γ_1, γ_2) of minimal length geodesics to level curves $\gamma_i \in \hat{G}(x, \chi_{u=c_i})$ where $c_1 \leq u(x) \leq c_2$ is called locally stable, if there exists $\varepsilon > 0$ such that for all $s < \varepsilon$ and $y = \tilde{\gamma}_1(s)$ or $y = \tilde{\gamma}_2(s)$, there is a minimal length geodesic pair (δ_1, δ_2) from y to the same level curves $\delta_i \in \hat{G}_{\Omega, g}(y, \chi_{u=c_i})$, such that*

$$\begin{aligned} \text{length}(\delta_1) &= \text{length}(\gamma_1) + \sigma s + o(s) \\ \text{length}(\delta_2) &= \text{length}(\gamma_2) - \sigma s + o(s) \end{aligned} \quad \text{where } \sigma = \begin{cases} +1 & \text{if } y = \tilde{\gamma}_2(s) \\ -1 & \text{if } y = \tilde{\gamma}_1(s) \end{cases} \quad (3.50)$$

Put another way, the geodesics are locally stable, if their total length does not change to the first order, when we move an infinitesimal distance along the geodesic.

Proposition 10. *Let u be a Lipschitz fixed point, whose geodesic pairs are locally stable. Consider any point x and any pair (γ_1, γ_2) of minimal length geodesics $\gamma_i \in \hat{G}_{\Omega, g}(x, \chi_{u=c_i})$ from the point x to the level curves. If the geodesic orientation follows the gradient at x , i.e. if*

$$\tilde{\gamma}_2'(0) \in \frac{D^+u}{|D^+u|}(x) \quad \text{and} \quad \tilde{\gamma}_1'(0) \in \frac{D^-u}{|D^-u|}(x)$$

then

$$|D^+u|(x) = |D^-u|(x) = \frac{c_2 - c_1}{L} \quad (3.51)$$

where $L = \text{length}(\gamma_1) + \text{length}(\gamma_2) = d(x, \chi_{u=c_1}) + d(x, \chi_{u=c_2})$ is the total length of the geodesic pair.

Proof. Since u is Lipschitz and its geodesics follow the gradient orientation at x , we can apply lemma 8 in the particular case of arc-length parameterization, to obtain

$$|D^+u|(x) = \left. \frac{d}{ds} u(\tilde{\gamma}_2(s)) \right|_{s=0}. \quad (3.52)$$

Now consider $x_h = \tilde{\gamma}_2(h)$, and the pair of geodesics $(\delta_{1,h}, \delta_{2,h})$ from x_h to the same level curves. Since u is a fixed point:

$$u(x_h) = \frac{c_1 \text{length}(\delta_{2,h}) + \text{length}(\delta_{1,h})c_2}{\text{length}(\delta_{2,h}) + \text{length}(\delta_{1,h})} \quad (3.53)$$

and since the geodesics are locally stable, for $h < \varepsilon$

$$\frac{u(x_h) - u(x_0)}{h} = \frac{c_2 - c_1 + o(1)}{L} \quad (3.54)$$

where $L = \text{length}(\gamma_2) + \text{length}(\gamma_1)$.

This shows that $|D^+u|(x) = \left. \frac{d}{ds} u(\tilde{\gamma}_2(s)) \right|_{s=0} = \lim_{h \rightarrow 0^+} \frac{u(x_h) - u(x_0)}{h} = \frac{c_2 - c_1}{L}$.

Similarly we show that $|D^-u|(x) = - \left. \frac{d}{ds} u(\tilde{\gamma}_1(s)) \right|_{s=0} = \lim_{h \rightarrow 0^+} \frac{u(x_0) - u(x_h)}{h} = \frac{c_2 - c_1}{L}$. \square

In the next lemma we establish sufficient regularity conditions for u , in order to meet the hypothesis of proposition 10, which is the basis of our main result. In fact we show global stability, while only local stability is needed, which suggests that there is still room for weakening the hypothesis and still keeping our main result.

Lemma 9. *Assume that a fixed point $u : \Omega \rightarrow \mathbb{R}$ is Lipschitz and C^1 with $|Du(x)| \neq 0$.⁸ Then for all $x \in \Omega$ the corresponding geodesic pairs are globally stable and*

$$|Du(x)| = |D^+u(x)| = |D^-u(x)| = \frac{c_2 - c_1}{L(x)} \quad (3.55)$$

where $L(x) = d(x, \chi_{u=c_1}) + d(x, \chi_{u=c_2})$ is the total length of a pair of minimal length geodesics from x .

⁸ The hypothesis that $|Du(x)| \neq 0$ is technical and can be possibly deduced from the fact that u is a fixed point whenever $c_1 < c_2$ (strict inequality), and the interior of Ω consists of a single connected component. The case $c_1 = c_2$ obviously leads to flat interpolators. So does a connected component of the interior of Ω whose boundary is entirely contained in either $\partial\Omega_1$ or $\partial\Omega_2$.

Proof. Applying proposition 9 we deduce that the geodesics γ follow the gradient direction $\frac{Du}{|Du|}(x)$ at almost every $x \in \Omega$. But since Du is continuous and non-zero, so is $\tilde{\gamma}$, so that actually $\tilde{\gamma}'(s) = \frac{Du}{|Du|}(x)$ for all $x \in \Omega$ and for all geodesics passing through $x = \tilde{\gamma}(s)$.

Thus the hypothesis of proposition 10 are satisfied everywhere, which allows us to deduce equation (3.55). \square

Finally we show that under the hypothesis of the lemma, fixed points are AMLE. Since these hypothesis do not allow us to take second order derivatives, we can either use viscosity solutions, or use Aronsson's characterization of AMLE as a minimal Lipschitz extension which is also minimal Lipschitz on each subdomain.

Theorem 13. *Under the hypothesis of lemma 9, if u is a fixed point then u is AMLE.*

Proof. First let's show that u is a minimal Lipschitz extension (MLE). This means that the Lipschitz constant $\text{Lip}(u, \Omega) = \sup_{x \in \Omega} |Du(x)|$ inside the domain, is the same as the Lipschitz constant $\text{Lip}(u, \partial\Omega)$ at the boundary which in our case is just $\text{Lip}(u, \partial\Omega) = \frac{c_2 - c_1}{d}$ where d is the geometric distance between the two level-curves $\partial\Omega_1 = \chi_{u=c_1}$ and $\partial\Omega_2 = \chi_{u=c_2}$. Actually it is sufficient to show that $\text{Lip}(u, \Omega) \leq \text{Lip}(u, \partial\Omega)$ because the inverse inequality is always true.

But this is very simple because from equation (3.55) we know that:

$$|Du(x)| \leq \frac{c_2 - c_1}{L(x)} \leq \frac{c_2 - c_1}{d} = \text{Lip}(u, \partial\Omega)$$

because no geodesic can be shorter than the distance d between the level lines. This shows that u is MLE.

Secondly we show that u is also absolutely minimizing (AMLE), and a sufficient condition is that for every subdomain $\Omega_1 \subseteq \Omega$ we also have $\text{Lip}(u, \Omega_1) \leq \text{Lip}(u, \partial\Omega_1)$.

Assume that this is not the case, so there are two points p and q inside Ω_1 such that $\frac{|u(p) - u(q)|}{|p - q|}$ is strictly larger than $\text{Lip}(u, \partial\Omega_1)$. This means that we can find x on the segment pq such that $|Du(x)| > \text{Lip}(u, \partial\Omega_1)$.⁹ Now consider a minimal length geodesic pair from x , with total length $L(x)$, which meets the boundary at x_1 and x_2 . According

⁹It is not necessary here that the segment pq be contained in the subdomain, since we can cut it into several pieces if necessary

to the lemma we have

$$\begin{aligned} |Du(x)| &= \frac{c_2 - c_1}{L(x)} = \frac{u(x_2) - u(x_1)}{L(x)} \\ &\leq \frac{u(x_2) - u(x_1)}{|x_2 - x_1|} \\ &\leq \frac{u(x_2) - u(x_1)}{d} = \text{Lip}(u, \partial\Omega_1) \end{aligned}$$

since the geodesic can't be shorter than the straight-line $|x_2 - x_1|$ which in turn can't be shorter than the distance d between the two parts of the boundary. But this contradicts the direct conclusion from our initial assumption that $|Du(x)| > \text{Lip}(u, \partial\Omega_1)$.

This shows that u is AMLE in the sense given by Aronsson. Hence by the uniqueness established by Jensen this means that u is the viscosity solution of $D^2u(Du, Du)$. \square

Possible generalization of the last lemma and theorem

The hypothesis of the last theorem are still too strong. It probably holds under slightly weaker assumptions. We can't accept for instance that u be non-differentiable along a whole interval of a geodesic, even if it is Lipschitz. In that case geodesics might be not even locally stable (this is the case at the edge of a pyramid for instance). However we can possibly accept that u be not C^1 at a few points, provided they are not dense. This is a first attempt in this direction, but the argumentation is still a bit confusing...

Lemma 10. *Assume that u is Lipschitz everywhere. In addition assume that within $\Omega \setminus P$ u is C^1 and $Du \neq 0$, where the set P is non-dense, i.e. for all $x \in \Omega$, there is $\varepsilon > 0$ such that $(B(x, \varepsilon) \setminus \{x\}) \cap P = \emptyset$. Then the corresponding geodesic pairs are globally stable for all $x \in \Omega$.*

In addition, for all $x \in \Omega$ we have

$$|Du|(x) \leq \sup_{y \in B(x, \varepsilon)} \frac{c_2 - c_1}{L(y)}. \quad (3.56)$$

where $|Du|(x) = \max\{|D^+u|(x), |D^-u|(x)\}$.

Proof. Consider a point $x \in \Omega$, and $\varepsilon > 0$ such that $(B(x, \varepsilon) \setminus \{x\}) \cap P = \emptyset$. Let (γ_1, γ_2) be a minimal length geodesic pair from x , and let $x_h = \tilde{\gamma}_2(h)$, for $h < \varepsilon$, and $(\delta_{1,h}, \delta_{2,h})$ a pair of minimal length geodesic curves. Applying proposition 9, for $s \in (0, h]$, we find that

$$\tilde{\gamma}_2'(s) = \frac{Du}{|Du|}(\gamma_2(s)) \quad \text{and} \quad \delta_{1,h}'(h-s) = -\frac{Du}{|Du|}(\delta_{1,s}(h-s)) \quad (3.57)$$

for almost every s in the interval. (we can write $\frac{Du}{|Du|}$ instead of $\frac{d^+}{|d^+|} \in \frac{D^+u}{|D^+u|}$ because u is C^1 with non-null gradient in this interval). But this means that $\tilde{\gamma}_2'(s)$ and $\delta_{1,h}'(h-s)$

are continuous for $s \in (0, h]$, so that actually equation (3.57) is valid for all $s \in (0, h]$. Since the curves are continuous, and they coincide at $s = h$, we conclude that for all $s \in [0, h]$

$$\tilde{\gamma}_2(s) = \delta_{1,h}(h - s).$$

So both pairs of geodesics coincide along the geodesic path between x and x_h . But then, the dynamic programming principle in equation (3.48) shows that they coincide along the whole geodesic path, *i.e.* they are globally stable.

Hence, the second part follows for $x \in \Omega \setminus P$ directly from proposition 10. For $x \in P$ observe that $|Du|(x) \leq \sup_{y \in B(x, \varepsilon)} |Du|(y) = \sup_{y \in B(x, \varepsilon)} \frac{c_2 - c_1}{L(y)}$. \square

The proof of theorem 13 follows in the same manner. We just substitute $|Du(x)|$ by the sup over a small neighborhood. However there might be an issue with the contradiction between the strict and non-strict inequalities that may become compatible with the sup.

A.4 Experiments

The following experiment illustrates the conclusions of the previous section.

We chose an interpolation domain where the AMLE solution is known to be C^1 but not C^2 . It consists of two level curves of a pyramid with square basis plus the summit.

In this case the convergence to the fixed point turned out to be very fast. At each iteration we computed the relative L^2 difference between the geodesic distance interpolator and AMLE, *i.e.* $d_k = \frac{\|u_{\text{amle}} - u_k\|}{\|u_{\text{amle}}\|}$ and we obtained the following values:

$$\begin{aligned} d_0 &= 0.0347 \\ d_1 &= 0.0164 \\ d_2 &= 0.0086 \\ d_3 &= 0.0065 \\ d_4 &= 0.0064 \end{aligned} \tag{3.58}$$

This means that in L^2 norm the fixed point comes after about four iterations very close to the AMLE and then it doesn't change significantly. From a qualitative point of view, we observe that the interpolator doesn't seem to be C^1 along the diagonals (as it can be observed in the level-lines display in figure 3.12), so that our theorem cannot be applied here. In fact, if it was C^1 along the diagonal, symmetry arguments show that the downwards geodesics from a point in the diagonal would follow the diagonal, and we would create a discontinuity. However, the irregularity that we obtain along the

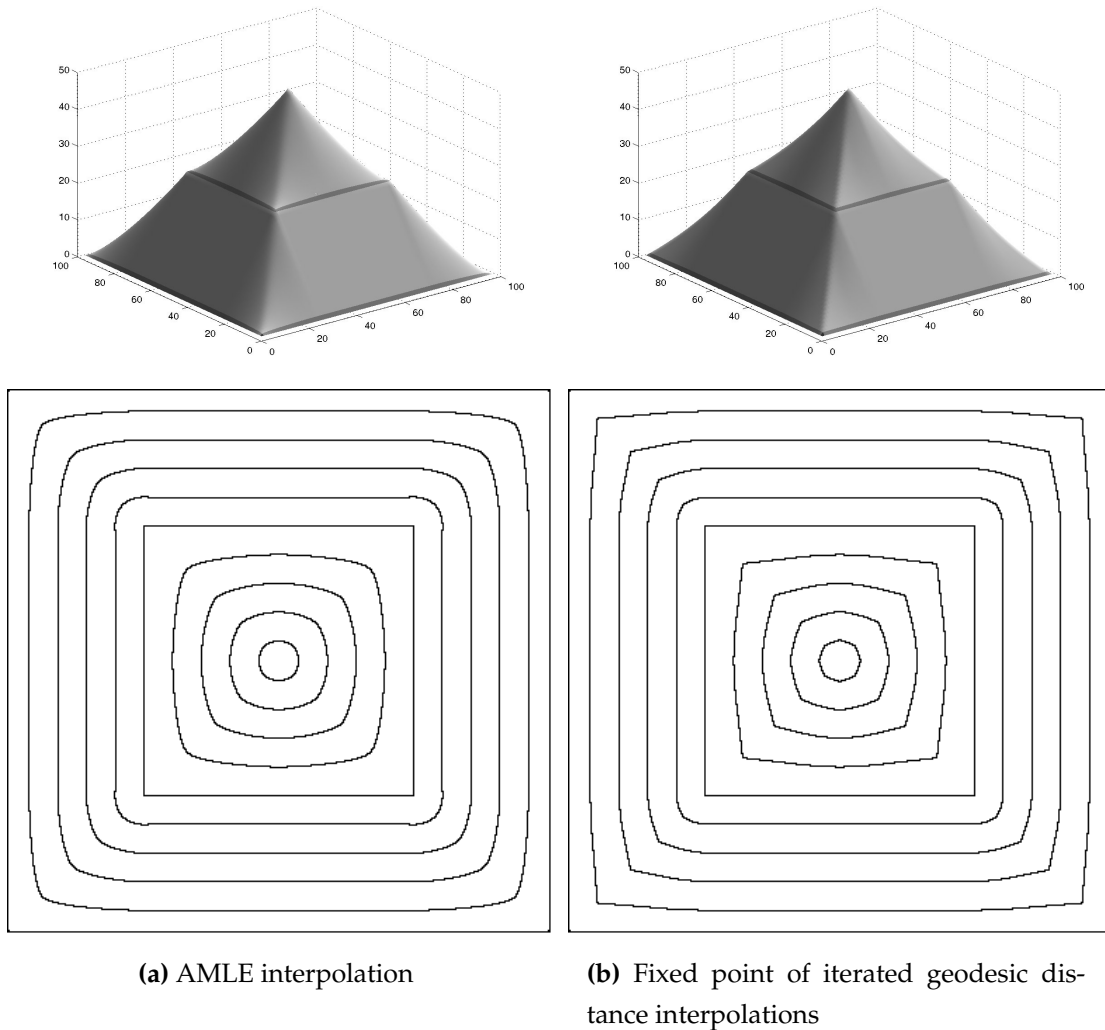


Figure 3.12: Comparison of AMLE and fixed point of geodesic distance interpolations

diagonal allows the downwards geodesics to escape from the diagonal, and we obtain a solution which is very close to AMLE.

This example suggests that the equivalence between AMLE and fixed points is not true in the general case. Fixed points seem to provide, nevertheless, a very fast initialization to a solution which is very close to AMLE, and can be later refined by a few iterations of the usual scheme to solve the AMLE.

B Numerical Analysis

In this section we discuss the numerical implementation of the two finite difference schemes we used for interpolating elevation data, namely the AMLE and the thin-

plate models. The computational complexity issues are discussed in section C, where the efficiency of these two methods is compared with Kriging and geodesic distance interpolations.

B.1 AMLE model using finite differences

We shall use the AMLE model studied above as the basic equation to interpolate data given on a set of curves and/or a set of points which may be irregularly sampled. In [42] a consistent finite-difference scheme was proposed for solving the AMLE. For the sake of completeness we reproduce here the main elements of their method. It consists of a particular discretization of the equation

$$D^2u \left(\frac{Du}{|Du|}, \frac{Du}{|Du|} \right) = 0$$

with boundary conditions

$$u(x) = u_0(x) \quad x \in \Omega.$$

It is easy to see that there is a relation between iterative methods for the solution of elliptic problems and time stepping finite difference methods for the solution of the corresponding parabolic problems. Because of that, we study the equation

$$\frac{\partial u}{\partial t} = D^2u \left(\frac{Du}{|Du|}, \frac{Du}{|Du|} \right),$$

with the following initial and boundary data

$$u(0, x) = u_0(x) \quad x \in \Omega \quad (3.59)$$

$$u(t, x) = \varphi(x) \quad (t, x) \in (0, +\infty) \times \partial\Omega. \quad (3.60)$$

Using an implicit Euler scheme we transform this evolution problem into a sequence of nonlinear elliptic problems. Thus, we may write the following implicit difference scheme in the image grid

$$u_{i,j}^{(n+1)} = u_{i,j}^{(n)} + \Delta t D^2 u_{i,j}^{(n+1)} \left(\frac{Du_{i,j}^{(n+1)}}{|Du_{i,j}^{(n+1)}|}, \frac{Du_{i,j}^{(n+1)}}{|Du_{i,j}^{(n+1)}|} \right) \quad (3.61)$$

$i, j = 1, \dots, N$. To solve the above nonlinear system we use a nonlinear over-relaxation method (NLOR). Writing the system as a set of $k = N^2$ algebraic equations, one for each unknown $u_{i,j}^{(n+1)}$ ($i, j = 1, \dots, N$),

$$f_p(x_1, x_2, \dots, x_k) = 0, \quad p = 1, 2, \dots, k,$$

the basic idea of NLOR is to introduce a relaxation factor ω and iteratively compute

$$x_i^{(n+1)} = x_i^{(n)} - \omega \frac{f_i(x_1^{(n+1)}, \dots, x_{i-1}^{(n+1)}, x_i^{(n)}, \dots, x_k^{(n)})}{f_{ii}(x_1^{(n+1)}, \dots, x_{i-1}^{(n+1)}, x_i^{(n)}, \dots, x_k^{(n)})}, \quad (3.62)$$

$$i = 1, 2, \dots, k$$

where $f_{ii} = \frac{\partial f_i}{\partial x_i}$. The convergence criterion can be shown to be the same as the over-relaxation method for linear systems, replacing the matrix by the Jacobian of the equations $f_p = 0$, and stability is guaranteed for values of the relaxation parameter $0 < \omega < 2$.

B.2 Thin-plate using finite differences

We assume that we are given a regular square grid $\Gamma = h\mathbb{Z}^2$, and that the values of u are known on $\tilde{\Omega}^c = \Gamma \cap \Omega^c$, and we want to obtain the values of u on $\tilde{\Omega} = \Gamma \cap \Omega$ satisfying equation (3.10) up to a certain tolerance.

To solve it, we could consider as in the case of the AMLE, the differential equation

$$\begin{aligned} \frac{\partial u}{\partial t}(x; t) &= -\Delta^2 u(x; t) \quad \text{for } x \in \Omega, t \geq 0 \\ u(x; 0) &= \varphi(x), \quad \text{for } x \in \partial\Omega \\ \frac{\partial}{\partial n} u(x; 0) &= \psi(x) \quad \text{for } x \in \partial\Omega \end{aligned} \quad (3.63)$$

(where the bilaplacian operator is taken with respect to the spatial x variable) and try to find the steady state when $t \rightarrow \infty$, for any initial condition.¹⁰ However, since in this case the differential operator is linear, it will be more convenient to directly discretize the original problem:

$$\begin{aligned} \Delta^2 u(x) &= 0 \quad \text{for } x \in \Omega \\ u(x) &= \varphi(x), \quad \text{for } x \in \partial\Omega \\ \frac{\partial u}{\partial n}(x) &= \psi(x) \quad \text{for } x \in \partial\Omega. \end{aligned} \quad (3.64)$$

The discretization of this equation leads to a linear system $L^2 U = W$ where $U_{i,j}$ is an approximation of the solution $u(hi, hj)$ on a regular grid of step h , L^2 is a second-order accurate discretization of the bilaplacian operator, which takes into account the

¹⁰This presents the advantage that this equation can be solved using explicit methods for which each iteration is inexpensive. But the stability condition in this case, for explicit methods is very tight and we are forced to take very small time-steps. Thus, we are forced to use explicit methods for each time step, which take as long to solve as the original, linear system.

$$\frac{V_{i,j}}{h^4} = \begin{cases} \begin{aligned} & U_{i,j-2} \\ & + 2U_{i-1,j-1} - 8U_{i,j-1} - 2U_{i+1,j-1} \\ & U_{i-2,j} - 8U_{i-1,j} + 20U_{i,j} - 8U_{i+1,j} + U_{i+2,j} \\ & + 2U_{i-1,j+1} - 8U_{i,j+1} - 2U_{i+1,j+1} \\ & U_{i,j+2} \end{aligned} & \text{if } (i,j) \in \tilde{\Omega} \\ U_{i,j} & \text{if } (i,j) \in (\tilde{\Omega} \oplus B) \setminus \tilde{\Omega} \end{cases} \quad (3.66)$$

boundary conditions, and W contains the boundary conditions. More precisely, if $V = L^2U$ then the vector V is defined as shown in Equation (3.66), and the right-hand side is

$$W_{i,j} = \begin{cases} 0 & \text{if } (i,j) \in \tilde{\Omega} \\ u_0(hi, hj) & \text{if } (i,j) \in (\tilde{\Omega} \oplus B) \setminus \tilde{\Omega} \end{cases} \quad (3.65)$$

where u_0 contains the known values of u outside of Ω . Note that L^2 operates on a vector U of values defined on the dilated domain $\tilde{\Omega} \oplus B$, and returns a vector V defined on the same domain. Here we denote by B a diamond-shaped structuring element matching the support of the discrete bilaplacian operator.

Now we need to find a convenient way to invert the operator L^2 in order to obtain U from W . Observe that L^2 is a very sparse linear operator, and numerical analysis provides many ways to limit the computational effort of inverting such an operator. Here we used a combination of biconjugate gradients with preconditioning, and multigrid. Multigrid is especially useful when interpolating on large domains, since it allows to quickly propagate the boundary data towards the interior.

Specifically the combined algorithm consists of the following steps:

1. Smooth the boundary values W by means of a 3×3 bilinear filter and resample on a coarser regular grid, to obtain a coarse vector W^1 containing four times less pixels than U .¹¹ Do the same for the indicator function of $\tilde{\Omega}$ and threshold it at $\theta \in (0, 1)$ to obtain the domain $\tilde{\Omega}^1$ on the coarse grid. Finally compute the discretization $L^{2,1}$ of the bilaplacian operator on the coarse grid.

¹¹ Here the filtering is done by taking into account only the boundary values, not the zeros contained in W when $x \in \Omega$. If, however, an initial approximation is available inside Ω , it can be used to filter W , and obtain a coarse initial approximation at the coarser level.

2. Iterate step 1 S times, obtaining $W^S, \tilde{\Omega}^S$ and $L^{2,S}$ which are coarse representations of W and $\tilde{\Omega}$ at scale level S .
3. Solve $L^{2,S}U^S = W^S$ by a conjugate gradient iteration, by using an incomplete LU factorization as a preconditioner, and taking any arbitrary initial condition.
4. Project U^S to the finer grid $S - 1$ by bilinear interpolation, and use this value as an initial condition to solve $L^{2,S-1}U^{S-1} = W^{S-1}$ by conjugate gradients. Here we are already quite close to the solution so only a few iterations should be needed. Hence, it is better not to use preconditioning, because its calculation is too expensive, compared to the savings.
5. Iterate step 4, S times, until we obtain $U = U^0$ on the original grid.

The optimal number of steps in the multigrid algorithm depends on the size of the region to be interpolated. We typically obtained the best results with two or three dyadic steps. We also obtained fast convergence in some cases without any multigrid step, just solving the preconditioned biconjugate gradient algorithm at the finest level. In certain cases, where the biconjugate gradient method converged too slow we switched to a successive over-relaxation (SOR) method for solving the system.

One final remark should be made regarding the boundary conditions. The way we discretized the bilaplacian operator ensures that the condition $u|_{\partial\Omega} = \varphi$ is satisfied (at least in a discrete sense, up to the approximation level of our grid). The second boundary condition, however, is more subtle. The following argument intends to show that imposing u on a dilated boundary $\tilde{\Omega} \oplus B - \tilde{\Omega}$ (where B is a structuring element of radius $2h$), actually corresponds to imposing $\frac{\partial u}{\partial n}$ on the boundary.

Consider, for simplicity the one dimensional case. Assume that $\Omega = (0, N)$ where $N > 5$. Then, according to our discretization of the bilaplacian operator (in 1D) the solution should satisfy

$$\begin{aligned} \Delta^2 u(h)h^4 + o(h^6) = \\ u(-h) - 4u(0) + 6u(h) - 4u(2h) + u(3h) = 0. \end{aligned} \tag{3.67}$$

The boundary condition $u(0)$ is directly imposed by the method. In order to find out if the interpolator meets the boundary condition $u'(0)$, we compare the given value $u'(l_h)$ with the interpolated value $u'(r_h)$ for two sequences $l_h \in \Omega^c$ and $r_h \in \Omega$ both converging to the boundary value 0. The boundary condition is met if $\lim_{h \rightarrow 0} u'(l_h) - u'(r_h) = u(0^-) - u(0^+) = 0$. Consider for instance $l_h = -0.5h$ and $r_h = 2.5h$. The

discrete approximation of the difference $u'(-0.5h) - u'(2.5h)$, can be written, using the previous equation as

$$\begin{aligned}
 \frac{h}{2}(u'(-0.5h) - u'(2.5h)) + o(h^2) &= \\
 &= -u(-h) + u(0) + u(2h) - u(3h) \\
 &= -3u(0) + 6u(h) - 3u(2h) \\
 &= 3\frac{h}{2}(u'(0.5h) - u'(1.5h)) + o(h^2).
 \end{aligned} \tag{3.68}$$

We can conclude that $u'(-0.5h) = 3u'(0.5h) - 3u'(1.5h) + u'(2.5h) + o(h)$. Finally let $h \rightarrow 0$ to conclude that $u'(0^-) = u'(0^+)$, which means that asymptotically, for a sufficiently fine grid, the derivative of the interpolator at $\partial\Omega$ (i.e. $u(0^+)$) matches the boundary condition $u'(0^-)$.

C Algorithmic complexity

In this section we briefly discuss the complexity of the algorithms we used for computing the different interpolation methods proposed. The computing times we needed on a SUN Ultra-SPARC processor at 336 MHz to run the examples in table 3.1 are given in table 3.2. These figures should be taken with care when comparing methods, because our implementations of these algorithms are not optimal in all cases, and they are implemented in different languages (C for AMLE and Matlab for thin-plate and Kriging). In the following paragraphs we give an indication of the computational complexity as well as some possible improvements in their implementations. This should be more useful to compare the performance of the different methods.

In all cases we assume that the region Ω to be interpolated consists of N grid-points, and that its boundary $\partial\Omega$ is composed of L grid-points. Furthermore, we call M the number of interpolation constraints, which in the case of AMLE and distance transforms is $M = L$ and in the case of thin-plate and Kriging is usually in the range $M \in [L, 2L]$. This is due to the fact that in order to impose slopes on the boundary, what we do in practice is impose the values on a dilated boundary, which is normally composed of about $2L$ points, except if data is not available on the dilated boundary as in the case of contour-line interpolation.

C.1 AMLE using finite differences

The numerical method in [42] and outlined in section B.1 consists of evolving the PDE $\frac{\partial u}{\partial t} = D^2u \left(\frac{Du}{|Du|}, \frac{Du}{|Du|} \right)$ with the same boundary condition $u(x; t) = \varphi(x)$, $(x, t) \in \partial\Omega \times (0, +\infty)$ and any initial condition $u(x; 0) = u_0(x)$, $x \in \Omega$ until a steady state is found (for $t \rightarrow \infty$), which satisfies the AMLE equation $D^2u(Du, Du) = 0$ up to a certain tolerance.

Each time step is discretized in an implicit manner, using one iteration of non-linear over-relaxation to solve it. The global complexity is then nNC flops, where n is the number of iterations needed to reach the steady-state, and $C \approx 50$ is the constant number of floating point operations per point and iteration. In our examples, $N \approx 200$ and we need about $n = 100$ iterations to achieve a precision of 2 decimal digits, and about $n = 200$ iterations for a precision of 4 digits. Total running times for these examples are 0.2 and 0.3 seconds respectively on a Sun SPARC workstation.

The number of iterations is highly dependent though on the size of the maximal gap in the interpolation domain. In a bigger example with $N \approx 100,000$ we need about $n = 3000$ iterations (10 minutes) for a precision of 2 decimal digits.

C.2 AMLE through iterated geodesic distance transforms

As we observed in section 3.1, in the special case of contour-line interpolation, the AMLE can be obtained (under certain hypotheses) by iterated weighted geodesic distance transform interpolations. The complexity of one iteration is CN flops, where C is a small constant, if we compute the weighted geodesic distances by a propagation algorithm like the one in [164], which visits each pixel just a few number of times. The total complexity of AMLE through geodesic distances is then nCN flops, but the number of iterations n is here much smaller than in the case of finite differences, typically two or three.

The iteration is not ensured to converge to the AMLE, but this can be used as a good initial approximation for the finite difference scheme, which would then take just a few iterations to converge to the steady state.

C.3 Thin-plate and Kriging through radial basis functions

Kriging and thin-plate interpolations can be both computed, as explained in section 3.3, by

1. Solving equation (3.19) in order to find the Kriging coefficient vectors \mathbf{a} and \mathbf{b} . This takes $\frac{1}{3}(M + N_k)^3$ flops in the case of k -IRF, where $N_k = 3$ for the common case of $k = 1$ which is used for $\alpha \in [1, 4)$, and then:
2. Using equation (3.18) to find the value of the interpolator at each point, which takes $(M + N_k)N$ flops.

So the total complexity is $\frac{1}{3}(M + N_k)^3 + (M + N_k)N$ flops, *i.e.* $\frac{1}{3}(M + 3)^3 + (M + 3)N$ in the particular case of $k = 1$, which includes the thin-plate. This is much faster than AMLE when M is small with respect to N (a roughly circular interpolation domain for instance), but much slower when M is large, which happens *e.g.* when the interpolation domain has a very irregular boundary.

C.4 Thin-plate through finite differences

In the special case of $\alpha = 2$ Kriging coincides with the thin-plate model and we can use a finite difference method to solve it. The spatial discretization of the biharmonic equation with boundary conditions (3.10) leads to a sparse (block-banded) $(N + M) \times N$ linear system with 21 bands (one for each point in the 5×5 stencil used to discretize the bilaplacian operator). To solve this system we used either a multi-grid approach, or a preconditioned conjugate-gradient method, with an incomplete LU factorization as a preconditioner. The total complexity is then $nm(N + M)$, where n is the number of iterations of the conjugate gradient algorithm and m is the number of non-zero elements per row in the incomplete LU factorization. In our experiments $m \approx 150$ and $n \ll (N + M)$, typically $n \approx 50$ for a four digit precision. Thus, for large values of M this method performs much better than the radial basis function method derived from Kriging.

Part III

Feature Detection

Chapter 4

Vanishing Point detection without any a priori information *

Abstract

Vanishing points in digital images result from the projection by a pin-hole camera of a set of parallel lines in 3D. Most of the proposed computational methods for detection of vanishing points are forced to rely heavily on additional properties (like orthogonality or coplanarity and equal distance) of the underlying 3D lines or knowledge of the camera calibration parameters in order to avoid spurious responses.

In this chapter we show that in many images of man-made environments vanishing points alone (without any additional assumption on the underlying 3D lines beyond pure parallelism) are meaningful Gestalts, in the sense of Helmholtz's perception principle, recently proposed for computer vision by Desolneux et al. [62, 64]. This leads to a vanishing point detector with a low false alarms rate and a high precision level, which doesn't rely on any a priori information on the image or calibration parameters, and doesn't require any parameter tuning.

The key differences of this detector with respect to previous approaches are the following: (i) The line segments are themselves detected with almost-zero false alarms rate, by a refinement of the method proposed by Desolneux et al. [61], also based on Helmholtz's principle; (ii) The threshold to determine a meaningful vanishing point from a large vote in the generalized Hough transform is computed in such a way that it guarantees a low false alarms rate; (iii) Finally a Minimum Description Length (MDL) criterion is used to further restrict the number of spurious vanishing points.

*A shorter version of this chapter was submitted jointly with Agnès Desolneux and Sébastien Vamech, and accepted for publication in an upcoming special issue on "Perceptual Organization in Computer Vision" of IEEE Transactions of Pattern Analysis and Machine Intelligence [15]. A preliminary version was published as a CMLA preprint [14], and in the proceedings of a special meeting at the Mathematisches Forschungsinstitut Oberwolfach [12], and presented at the Image Analysis seminar of the Jacques-Louis Lions Laboratory, Université de Paris 6.

1 Introduction

Sets of parallel lines in 3D space are projected into a 2D image obtained with a pin-hole camera to a set of concurrent lines. The meeting point of these lines in the image plane, is called a *vanishing point* [70], and may eventually belong to the line at infinity of the image plane in the case of 3D lines parallel to the image plane. Even though concurrence in the image plane doesn't necessarily imply parallelism in 3D (it only implies that all 3D lines intersect the line defined by the focal point and the vanishing point), the counterexamples for this implication are extremely rare in real images, and the problem of grouping sets of parallel lines in 3D and their corresponding vanishing points in the image plane is reduced to finding significant sets of concurrent lines in the image plane.

What are vanishing points needed for? The usefulness of precise measurements of vanishing points, among other geometric primitives, was demonstrated in many different frameworks [54, 70, 118, 154].

A common situation in architectural environments is to find a set of three orthogonal dominant orientations. If the corresponding vanishing points are detected in the image, they provide three independent constraints on the 5 internal calibration parameters of the camera [118]. More importantly, it is very common to have cameras with zero skew and aspect ratio equal to 1 (a natural camera). In this case the internal parameters are reduced to 3 (the focal length, and the position of the principal point (the orthogonal projection of the focal point into the image plane), and the camera can be calibrated from a single image of a building, for instance, with the only assumption that the walls and the floor are orthogonal to each other [118].

Even more impressive is the result described in [54], where it is shown that (without any knowledge of camera calibration or position) the ratios of lengths between two parallel segments in 3D can be computed from the lengths of the imaged segments, if we only know the vanishing line of a plane containing one of the endpoints of both segments, and the vanishing point of the two parallel segments. A typical application of this result is to measure the height of objects standing (vertical) on the floor, relative to a reference object standing (vertical) on the floor parallel to them.¹ The only points that are required are: two horizontal vanishing points, a vertical vanishing points, and

¹ Note that here we used the word "vertical" for clarity and to fix ideas, but actually no orthogonality relationship is required, only 3D parallelism.

the endpoints of the target and reference segments. Then the computations only involve measuring ratios of lengths of segments defined by these points.

Finally when using vanishing points in conjunction with other properties such as orthogonality, common segments between planes, and a few length ratios, and/or knowledge that the camera is a natural camera), full 3D reconstructions are possible from a single view, even of Renaissance paintings which carefully respected perspective laws [118].

When several views are available, vanishing point correspondences can be helpful in determining the epipolar geometry or fundamental matrix [70]. Alternatively if the intrinsic parameters are known, vanishing points can be used to find some information about extrinsic parameters, *i.e.* the relative position of the two cameras. In cite [23] for instance, vanishing points were used to decouple the relative rotation between the two cameras, from the relative translation to which vanishing points are insensitive.

Available methods for vanishing point detection. Since the seminal work of Barnard [169], automated computational methods for vanishing points detection in digital images have been based on some variation of the Hough transform in a conveniently quantized Gaussian sphere. Several refinements of these techniques followed, but most recent works suggest that this simple technique often leads to spurious vanishing points[160]. In order to eliminate these false alarms most authors considered some kind of joint Gestalt, which adds some other property to 3D parallelism like coplanarity and equal distance between lines [156] or orthogonality between the three main 3D directions [120, 146, 160]. In addition, knowledge of the intrinsic camera calibration parameters is commonly assumed [23, 120] by these methods, or they are designed mostly for omnidirectional images [23]. To the best of our knowledge, the question of reliably determining whether an image *actually contains* some vanishing points and its *number*, has not yet been addressed systematically.

The proposed approach. In this chapter we show that 3D parallelism alone is a significant Gestalt in many images of man-made environments and that it can be reliably and automatically detected with a low number of false alarms and a high precision level, without using any secondary property, or any a priori information on the image or calibration parameters, and without any parameter-tuning. We do not claim that secondary properties (like equal distance, or orthogonality) should not be used in any circumstance; this can be useful for some applications, and our technique could be eventually extended to these situations. But in many applications a pure vanishing

point detector is more useful, since it can be used to determine some calibration parameters of the camera (which are needed in other approaches relying on orthogonality for instance), or for other applications such as single view metrology as explained above. The key improvements with respect to previous vanishing point detectors are the following:

1. The primitives that are accumulated in (an equivalent of) the Gaussian sphere are line segments, which are themselves detected with an almost-zero false alarms rate, by a refinement of the method presented in [61].
2. Our criterion to determine a meaningful vanishing point from a large vote in the Gaussian sphere is deduced from the Helmholtz principle, thus producing a low number of false alarms, without need for threshold-tuning.
3. Finally a Minimum Description Length (MDL) criterion is used to further restrict the number of spurious vanishing points and to deal with the masking phenomenon.

The rest of the chapter is organized as follows. Section 2 gives some background on the Gestalt principles used in this chapter, and presents the refinements of the alignment detection algorithm, that were necessary to adapt it for vanishing point detection. Section 3 describes the vanishing point detection algorithm itself, and Section 4 presents the results of our experiments.

2 Detection of Line Segments

As a primitive for our vanishing point detector we use a set of line segments (edges) that have been obtained with a refinement of the method presented in [61]. In the following subsection 2.1 we briefly review the method and its philosophy, then we discuss the limitations that should be addressed in order to produce useful primitives for vanishing point detection and we describe our proposed refinements (subsection 2.2). The initial review of the line segment detector will also serve the reader as an introduction to vanishing point detection, since we follow exactly the same methodology.

2.1 Helmholtz principle and alignment detection

The general guiding principle that will be used throughout this chapter is due to Helmholtz and was recently proposed in [62, 64] as a general technique for feature detection:

Assume we are observing n independent objects, and that k out of them share a given common property. Then the grouping of these k objects is perceived as a Gestalt if the event $G =$ “at least k out of n objects share the property” is extremely rare in a random setting.

The random setting can be for instance the assumption that the observed property is independent for all the n objects with probability p of occurring. Hence the probability of observing the configuration G becomes the tail of the binomial distribution:

$$B(p, n, k) = \sum_{i=k}^n \binom{n}{i} p^i (1-p)^{n-i}. \quad (4.1)$$

The notion of “extremely rare” is formalized as the expected number of occurrences of the event being smaller than a small value ε . This expectation can be upper-bounded simply by $N_T B(p, n, k)$, where N_T is the number of times we test for an event like G in the image. Now if we actually observe a group G of k among n objects sharing the given property, the previous observations lead us to define the number of false alarms of the geometric configuration G as:

$$\text{NFA}(G) := N_T B(p, n, k). \quad (4.2)$$

We also say that G is an ε -meaningful event if $\text{NFA}(G) \leq \varepsilon$. These definitions ensure that the expected number of detections of an event like G under a random setting is lower than ε .

In the case of alignments [61], the objects we consider are pixels, the groups of objects we shall consider are sets of pixels forming discrete line segments, and the common property we shall test for is whether the gradient direction of these pixels coincide (up to a certain precision level to be specified) with the direction orthogonal to the line segment.

Now assume we are observing a segment A of length l , and that the gradient orientation can be measured with a precision $p\pi$. Then the probability that the gradient at a single pixel be aligned with that of the line segment is p , and the probability that at least k among the l independent pixels be aligned is $B(p, k, l)$. On the other hand the total number of tests on an $N \times N$ image is $N_T = N^2(N^2 - 1) \approx N^4$ which corresponds to the total number of possible oriented discrete line segments. So the number of false alarms for such an alignment A is defined as

$$\text{NFA}(A) := N^4 B(p, l, k), \quad (4.3)$$

and the alignment is ε -meaningful if $\text{NFA}(A) \leq \varepsilon$.

Highly significant segments may still be meaningful if we consider a subset or a superset of this segment, even if the superset contains no aligned points or if the subset contains only aligned points. To avoid this kind of “spurious responses” we will consider, among a family of nested ε -meaningful segments, only the most significant one, *i.e.* a segment A is *maximal meaningful* if it is meaningful and:

1. it does not contain any more meaningful segment, *i.e.* $\forall B, B \subseteq A \Rightarrow \text{NFA}(B) \geq \text{NFA}(A)$
2. it is not contained in any more meaningful segment, *i.e.* $\forall B, B \supseteq A \Rightarrow \text{NFA}(B) \geq \text{NFA}(A)$

2.2 Limitations of maximal meaningful alignments and proposed refinements

Multiple responses. The main limitation of the maximal meaningful alignments just described [61] is that we may obtain several candidates for each segment as it was shown in [13]. This is due to the fact that correctly sampled images are at least slightly blurred, which means that edges are a little bit thicker than one or two pixels. Hence many thin segments contained in the actual thick segment are still meaningful. Among all these segments we are interested in selecting a single one, namely the one which best estimates its position and orientation. We do so because multiple responses may lead to : (i) a larger localization error of the detected vanishing point (since only a few of these multiple responses are at the correct orientation while the others are slightly deviated from it but still meaningful); or (ii) accidental vanishing points (since a set of multiple responses might itself become an artificial vanishing point). The selection of the best candidate among multiple responses can be achieved by the following criterion:

MDL with wide segment. Once all maximal meaningful segments have been found, we start a competition between neighboring segments in order to determine the best candidate. The competition is based on the principle that each pixel belongs to a single segment, namely the most meaningful segment among all meaningful segments meeting that point. More precisely a pixel x is assigned to segment A iff $\forall B, (\text{NFA}(B) \leq \varepsilon) \wedge (x \in B \oplus D_r) \Rightarrow \text{NFA}(B) > \text{NFA}(A)$. (Note that segment B competes not only for its own pixels, but also for pixels in a slightly dilated segment $B \oplus D_r$, with D_r a disk of radius $r \approx 1$). Then we recompute $\text{NFA}(A)$ for all maximal meaningful

segments using equation (4.3), with the only modification that instead of the number k of aligned points in A is now lower, we use the number $k' \leq k$ of aligned points in A which in addition *have been assigned to the segment*. Therefore $\text{NFA}(A)$ increases, and if $\text{NFA}(A)$ is still smaller than ε the segment is called a *maximal MDL meaningful segment*. A similar criterion (with $r = 0$) was already considered in [61] in order to avoid oblique edges that may appear as a side effect of several parallel edges. With the modification we propose here ($r = 1$), the same criterion can be used to eliminate multiple almost parallel candidates of a blurred edge.

Canny points. In [13] we considered a second, more classical, criterion to avoid multiple responses for a single edge. It consisted of considering only alignments of Canny points, which are defined as those pixels where the Canny operator $D^2u(Du, Du)$ has a zero crossing. It produces similar results to the wide-segment MDL criterion, but it is more prone to eliminate important segments altogether as it can be observed in figure 4.1. On the other hand, it has the good property of allowing a significant acceleration of the algorithm. Currently we are looking for ways to use this criterion in conjunction with MDL, in order to accelerate the algorithm without missing any segments beyond multiple responses. For the results on vanishing point detection presented in this chapter we only used the wide segment MDL criterion.

Quantization of gradient orientation. In [58] and [59, chapter 3], it has been pointed out how the quantization of the gray-levels in an image translates into a very non-uniform distribution of the gradient orientation, leading in turn to a large number of wrongly detected segments at certain preferred orientations if this problem is not dealt with properly. In the same work two possible solutions are proposed. The first one consists of translating the image by half a pixel in both directions using Fourier interpolation. This transform actually converts the quantization noise and the corresponding noise in gradient orientation into quasi-white noise, and quasi-uniform noise respectively. The second solution consists of thresholding the gradient magnitude, and substituting the orientations of gradients below the threshold by uniform noise in $[0, 2\pi]$. The choice of the threshold on the gradient magnitude is given by

$$|Du| \geq \frac{q}{p\pi} \quad (4.4)$$

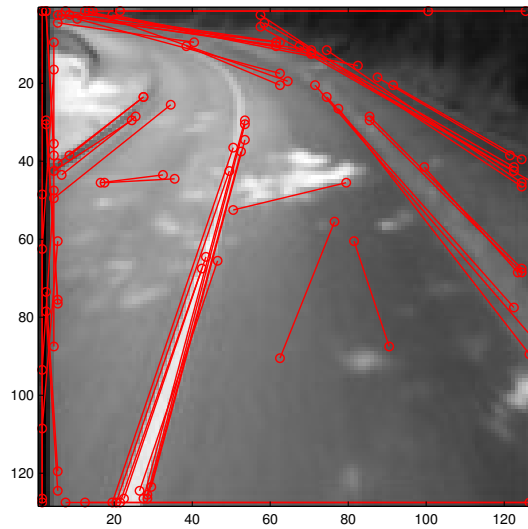
where q is the amplitude of the image quantization noise, and $p\pi$ is the angular precision for testing the alignment of gradients ($1/p$ different orientations). For smaller gradients, the orientation noise due to image quantization would deviate too much

from uniform, but for larger gradients the uniform approximation is acceptable. For the usual angular precision $p = \frac{1}{16}$, this means thresholding gradients smaller than 5 times the quantization step (usually 1 gray-level). But for finer precisions we will require a higher threshold (e.g. 20 gray-levels for $p = \frac{1}{64}$). Here we adopted this second solution (with a different gradient threshold at each precision level) because it has the advantage that it allows for sub-pixel accuracy in the location of the endpoints of detected line segments, whereas in the first solution we can only evaluate the gradient at half-pixels.


Parameters. It would seem that the method we just described depends on two parameters: the precision p and the false alarms threshold ε . However, we can say that the method is parameterless, for the following reasons. First we can fix both parameters once and for all at $\varepsilon = 1$ and $p = \frac{1}{16}$ and these values worked well in all our experiments. Secondly, the choice of ε has been shown not to be critical at all. In fact the results do not change significantly if we set, say $\varepsilon = 0.1$ or $\varepsilon = 10$, since the output of the algorithm varies rather like $\log \varepsilon$, i.e. very slowly. Finally, for the precision parameter p the chosen value is actually an optimal compromise: less precise values would require much too long segments to be meaningful, while more precise values would require edges which are very contrasted with respect to image noise to obtain a sufficient number of aligned points. In [13] we described a method to automatically select the optimal p for each segment. As a rule no segment is detected beyond precisions $p = 1/4$ and $p = 1/64$, while 8, 25, 45, 20, 3 percent of the segments chose precisions $p = \frac{1}{4}, \frac{1}{8}, \frac{1}{16}, \frac{1}{32}, \frac{1}{64}$ respectively. In addition most of the segments that preferred precisions other than $p = \frac{1}{16}$, are also detected at $p = \frac{1}{16}$, only with a larger NFA. For these reasons in our experiments we fixed always $p = \frac{1}{16}$. Images taken at extreme contrast or noisy conditions or very aliased images may require the adaptive method that automatically selects the best p .

3 Detection of Vanishing Points


As in the case of alignments we shall define a meaningful vanishing point in terms of the Helmholtz principle. Our objects in this case will be all the meaningful segments obtained by the method we described in the previous section. The common property we shall seek among these segments is a common point v_∞ met by all their supporting



(a) Maximal meaningful segments at precision $p = 1/16$ (63)



(b) Multi-precision Maximal MDL meaningful segments (14)



(c) Multi-precision Maximal Canny MDL meaningful segments (11)

Figure 4.1: Comparison between: (a) the method proposed in [61] ; and the adaptations we proposed to adapt them for vanishing point detection: (b) multi-precision and MDL, (c) multi-precision, Canny and MDL.

lines². Due to measurement errors, we shall never find a large number of segments intersecting in a single point v_∞ ; we shall rather find a family of lines intersecting a more or less small subset V of the image plane, which we shall call *vanishing region*. To consider all possibilities we need to consider a *finite* family of such regions $\{V_j\}_{j=1}^M$, such that it covers the whole (*infinite*) image plane, *i.e.*

$$\bigcup_{i=j}^M V_j = \mathbb{P}_2. \quad (4.5)$$

In [170], an intelligent such partition is proposed whereas most works use a partition of the image plane such that the projection of each vanishing region on the Gaussian sphere has a quasi-constant area [120, 146, 160, 169]. This partition has the advantage that it assigns the same precision to all 3D orientations, but it requires knowledge of the internal camera calibration parameters. However a practical application of the Helmholtz principle (Sections 3.1 and 3.2) leads to a different partition of the image plane into vanishing regions (to be introduced in Section 3.3), which shares some qualitative properties with the common Gaussian sphere partition. Further refinements of the proposed detection method are presented in Sections 3.4 and 3.5.

3.1 Meaningful vanishing regions

Here we adapt the Helmholtz principle to the case where the objects that are observed are the supporting lines l_1, l_2, \dots, l_N of the N line segments found in the previous section. The common property we shall test for, is whether a group of k such lines intersects one of the vanishing regions V_j . Figure 4.2(a) illustrates this construction. Under the assumption that all lines are independent with the same distribution, the probability of such an event is $B(p_j, N, k)$, where p_j is the probability that a line meets the vanishing region V_j . Moreover, since the M regions V_j are chosen to sample all possible vanishing regions, we make $N_T = M$ such tests. Thus the number of false alarms for a vanishing region V_j can be defined as:

$$\text{NFA}(V_j) := MB(p_j, N, k), \quad (4.6)$$

and as usual, the vanishing region is ε -meaningful if k is sufficiently large to have $\text{NFA}(V_j) \leq \varepsilon$.

²If the segments are pin-hole projections of 3D lines with a common orientation d_∞ , then this point v_∞ is the pin-hole projection of the 3D orientation d_∞

In order to actually find the value of NFA and the value of the minimal value $k(j, \varepsilon)$ of k such that V_j becomes meaningful, we need to know the probabilities p_j . This is the subject of the next section.

3.2 Probability of a line meeting a vanishing region.


Up to this point the analysis is formally almost equivalent to the case of meaningful alignments, stressing the duality of the events “ n lines meet at a point” and “ n points belong to a single line”. The threshold $k(j, \varepsilon)$ is computed in the same manner and from the same binomial tail as in the previous case. The only thing that changed is the interpretation of the parameters: the total number of segments N plays here the role of the length of the segment in the case of alignments, and the number of events M represents the total number of possible vanishing regions (instead of N^4 , the total number of possible segments). The specific geometry of the vanishing point problem comes into play only at this point when computing the probability p_j of a random line meeting (the image and) a vanishing region V_j .

Gratefully this geometric probability problem has been very elegantly solved in [155] who gives a closed form formula in terms of the internal and external perimeters of both regions. Here we state the main result from integral geometry that is needed to compute p_j . We refer the reader to appendix A for a summary of the main ideas of the proof, and to the treaty by Santaló [155] for a complete development of the theory leading to this and other interesting results.

First the polar coordinates parameterization for a random line on the plane G is considered (see figure 4.3(a)):


$$G(\rho, \theta) = \{(x, y) \in \mathbb{R}^2 : x \cos \theta + y \sin \theta = \rho\} \quad (4.7)$$

and it is shown from symmetry arguments that the only translation and rotation invariant measure for sets of lines is $dG = d\rho d\theta$ (up to an irrelevant multiplicative constant). It is important to be careful when choosing this measure, because it is not self-evident, and we could be easily misled by wrong arguments as it is illustrated by the Bertrand paradox in figure 4.4. We encourage the reader to look at the figure and find the right answer by himself. Evidently each solution corresponds to a different measure for the set of random lines, which corresponds to a different concept of “symmetry”. But only one of the answers corresponds to the rotation and translation invariant measure for sets of lines.



N segments

(a) Vanishing region V and a set of lines meeting it



(b) Interior vanishing region

(c) Exterior vanishing region

Figure 4.2: *Meaningful Vanishing regions.* (a) The problem consists of estimating the expected number of occurrences of the event “at least k out of N lines meet a vanishing region V_j ”, given that the lines meet (are visible in) the image domain Ω . In order to compute the associated probabilities $p = P[G \text{ meets } V_j \mid G \text{ meets } \Omega]$ we distinguish two cases: (b) Interior vanishing regions $V_j \subseteq \Omega$: in this case we have $p = \text{Per}(V_j) / \text{Per}(\Omega)$; (c) Exterior vanishing regions $V_j \cap \Omega = \emptyset$: in this case we have $p = (L_i - L_e) / \text{Per}(\Omega)$.




Figure 4.3: (a) With this parameterization the measure $dG = d\rho d\theta$ becomes translation and rotation invariant. (b) The support function $p(\theta)$ defined in equation (4.17) uniquely determines the convex set K and plays a central role when computing the probability of a line meeting a convex set K .

Once we chose this measure we can formulate the main result from integral geometry that we shall use here.

Proposition 11. Consider two convex sets K_1 and K_2 of the plane. Then the measure of all lines meeting both sets is:

$$\mu[G \cap K_1 \neq \emptyset \text{ and } G \cap K_2 \neq \emptyset] = \begin{cases} L_2 = \text{Per}(K_2) & \text{if } K_1 \subseteq K_2 \\ L_i - L_e & \text{if } K_1 \cap K_2 = \emptyset \\ \text{Per}(K_1) + \text{Per}(K_2) - L_e & \text{otherwise} \end{cases} \quad (4.8)$$

where the external perimeter L_e is the perimeter of the convex hull of K_1 and K_2 and the internal perimeter L_i is the length of the “internal envelope” of both sets, which is composed of the internal bitangents to K_1 and K_2 and parts of their perimeters. Figure 4.2(c) illustrates this construction for $K_1 = \Omega$ and $K_2 = V_j$.

This result can be directly applied to our problem of determining p_j in the case where the vanishing region $V_j \subseteq \Omega$ is contained in the (convex) image domain Ω . Since we can only observe line segments that intersect the image domain, the probability we




Figure 4.4: “Bertrand paradox”: Given a random line G which meets the circumscribed circle C of a regular triangle of side a , what’s the probability that the length of the cord σ be larger than a ? (a) Formulates the problem; (b) The event $\sigma > a$ is equivalent to the event “ G meets the inner circle C' ” with radius $r/2$. Since $\text{area}(C') = \text{area}(C)/4$ we conclude that $p = 1/4$. (c) Without loss of generality fix one side of the cord at the point A . Then $\sigma > a$ iff the other side of the cord B , is within an arc of length $1/3$ of the whole perimeter, and we conclude that $p = 1/3$; (d) Without loss of generality assume that the cord is horizontal. Then those meeting the condition are within the gray zone, whose width is half the diameter of the circle, so $p = 1/2$.

are interested in is actually:

$$\begin{aligned} p_j &= P[G \cap V_j \neq \emptyset \mid G \cap \Omega \neq \emptyset] = \frac{\mu[G \cap V_j \neq \emptyset \text{ and } G \cap \Omega \neq \emptyset]}{\mu[G \cap \Omega \neq \emptyset]} \\ &= \frac{\mu[G \cap V_j \neq \emptyset]}{\mu[G \cap \Omega \neq \emptyset]} = \frac{\text{Per}(V_j)}{\text{Per}(\Omega)}. \end{aligned} \quad (4.9)$$

For vanishing regions $V_j \cap \Omega = \emptyset$ external to the image domain, we just apply the second case of equation (4.8) and the probability becomes

$$\begin{aligned} p_j &= P[G \cap V_j \neq \emptyset \mid G \cap \Omega \neq \emptyset] = \frac{\mu[G \cap V_j \neq \emptyset \text{ and } G \cap \Omega \neq \emptyset]}{\mu[G \cap \Omega \neq \emptyset]} \\ &= \frac{L_i - L_e}{\text{Per}(\Omega)}. \end{aligned} \quad (4.10)$$

Note that the intermediate case where there is an intersection but no inclusion is treated as this second case with $L_i = \text{Per}(K_1) + \text{Per}(K_2)$, which is true in the limiting case when K_2 is tangent but exterior to K_1 . In the other limiting case, when K_2 is tangent but interior to K_1 , we still have $L_i = \text{Per}(K_1) + \text{Per}(K_2)$, but $L_e = \text{Per}(K_1)$, so $L_i - L_e = \text{Per}(K_2)$ and we are back in the first case.

3.3 Partition of the image plane into vanishing regions

In this section we address the problem of choosing a convenient partition of the image plane into vanishing regions. For this purpose we use the following criteria:

Equal probability. We try to build a partition such that the probability $p_j = P[G \cap V_j \neq \emptyset]$ that a random line G of the image meets a vanishing region V_j is constant for all regions. Without this equiprobability condition, certain vanishing regions would require many more meeting lines to become meaningful than others, *i.e.* they would not be *equally detectable*, which is not desirable.³ We can easily deduce from the results of the previous section that this equiprobability condition implies that the size of V_j increases dramatically with its distance from the image, which agrees with the fact that the *localization error* of a vanishing point increases with its distance from the image. Thus, with the equiprobability condition, we obtain the localization error of the vanishing points as a consequence of their detectability.

³For instance the partition into regions whose projection into the Gaussian sphere has constant area, does not necessarily satisfy this equal probability condition. This was observed by [120] in the case of uniformly distributed 3D lines. In this case, lines almost parallel to the image plane become much less probable than lines which are almost orthogonal. Despite the correction proposed in [120], this still leads to problems in the detection of vanishing points when the perspective effect is very low (distant vanishing points, or lines almost parallel to the image plane), as observed by [160].




Figure 4.5: *Uncertainty cone of a line segment [160]* . When a line segment of length l has been detected in a digital image, we can consider that the position of its endpoints has an uncertainty of about half a pixel. Hence the supporting line of this segment lies within an “uncertainty cone” centered at the baricenter of the segment, and with angle $d\theta = \sin^{-1} \frac{1}{l}$.

Angular precision. The size and shape of the vanishing regions should be in accordance with the angular precision of the detected line segments. In [160] the author addresses this problem by considering a localization error of 1 pixel at the ends of the segment, so the precision of the segment’s orientation is $d\theta = \sin^{-1} \frac{1}{l}$, where l is the length of the segment. The supporting line of the segment should be rather considered as a “cone” with angle $d\theta$ (see Figure 4.5). When such a cone intersects a vanishing region the corresponding accumulator is updated by a value proportional to the angular fraction of the cone covered by the vanishing region. This fraction becomes 1 if the vanishing region is larger than the width of the cone. Using Shufelt’s concept within our framework would require complex geometric probability calculations for which we do not have a closed form formula yet. So we use a similar concept which resembles thresholding by not considering fractional intersections. If an alignment has a sufficiently precise orientation, for its uncertainty cone to be completely contained in a vanishing region, then we count the intersection of the supporting line and the vanishing region. Otherwise, the intersection is uncertain, and we consider that the segment does not meet the vanishing region. This way we still count intersections of convex sets with lines (for which we know how to compute probabilities), instead of intersections of convex sets and cones (for which we don’t have a closed form expression). Despite this modification we still consider the angular precision of detected segments in a way similar to Shufelt’s work. This leads us to construct the vanishing regions in such a way that their size is comparable to the width of the corresponding vanishing cones.




Figure 4.6: Construction of the exterior vanishing regions V_j , and computation of the corresponding probability $(L_i - L_e)/\text{Per}(\Omega)$ that a random line meeting the image domain Ω also meets V_j . Observe that the correspondence between the angular sizes of internal and external tiles implies that $p_e = p_i = \frac{4\sin(d\theta)}{\pi}$. To compute the probability p_e of a line meeting an external tile V_e observe that $L_i - L_e$ is composed of two arcs of circle of angle α , the upper and lower sides of V_j of length $q' - q = \frac{d' - d}{\cos\theta}$, and two line segments of length $l_i - l_e$. Hence $L_i - L_e = 2(R\alpha + q' - q + l_i - l_e)$. Furthermore the right triangles of angles β and β' imply that $l_i = R \tan \beta$ and $l_e = R \tan \beta'$. Finally, since $\alpha + \beta = 2\theta + \beta'$, and $\text{Per}(\Omega) = 2R\pi$, substituting all equations into (4.10) we obtain $p_e = \frac{1}{\pi} \left(2\theta + \left[\beta + \frac{1}{\cos\beta} - \tan\beta \right]_{\beta=\arccos(R \cos \theta/d)}^{\beta=\arccos(R \cos \theta/d')}$.




Figure 4.7: Tilings of the image plane by equal probability vanishing regions for different angular precision levels. Only the exterior tiles are shown, except the last ring of "infinite" tiles, whose probability may be smaller. The interior tiles form a regular square partition of the circle, with squares of size comparable to the "inner-most ring" of exterior tiles. Observe how the size of the tiles increases for the more distant tiles. The axes represent distances relative to the radius of the circular image domain. Note that for higher angular precisions the image has been zoomed out, to allow the visualization of the more distant tiles.

Geometric construction. Now we shall construct a partition of the plane into vanishing regions that closely satisfies both criteria above. The partition is composed of two families of vanishing regions $\{V_j^{(i)}\}$ and $\{V_j^{(e)}\}$. The first (“interior”) one consists of regions entirely contained in the image domain Ω , and the second (“exterior”) one consists of regions outside the image domain.

For simplicity, we shall approximate the image domain by its circumscribed circle, and consider the image domain Ω as a circle of radius $R = N/\sqrt{2}$. In order to meet the angular precision requirement, all exterior regions V will be portions of sectors of angle $d\theta$ lying between distances d and d' from the image center O . Figure 4.7 illustrates this construction and figure 4.6 the trigonometric calculation of the probability that a random line meeting the image domain Ω does also meet V . In the case of exterior tiles of angular precision $d\theta$ at distances d and d') this probability becomes

$$p_e(d, d') = \frac{L_i - L_e}{\text{Per}(\Omega)} = \frac{1}{\pi} \left(2d\theta + \left[\beta + \frac{1}{\cos \beta} - \tan \beta \right]_{\beta=\arccos(R \cos(d\theta)/d)}^{\beta=\arccos(R \cos(d\theta)/d')} \right). \quad (4.11)$$

Note that it will be occasionally more handy to think of p_e as a function of the angles $\beta = \arccos(R \cos(d\theta)/d)$ and $\beta' = \arccos(R \cos(d\theta)/d')$ instead of the distances d and d' .

Concerning the interior regions, we chose a simple tiling of the circle Ω with square tiles. The side of each square is chosen to be equal to the side of the exterior tiles closest to the image domain, *i.e.* $2R \sin(d\theta)$. The perimeter of the interior regions is therefore equal to $8R \sin(d\theta)$, and the probability that a line meets an interior vanishing region is:

$$p_i = \frac{\text{Per}(V)}{\text{Per}(\Omega)} = \frac{4 \sin(d\theta)}{\pi} \quad (4.12)$$

This ensures that all interior regions have the same probability, and that their size is in accordance with the coarsest angular precision $d\theta$ of the line segments. Now we have to choose the values of d and d' to ensure that all exterior regions have the same probability $p_e = p_i$. To do so, we start with the first ring of exterior regions setting $d_1 = R$, and we choose d'_1 by solving the equation $p_e(d_1, d'_1) = p_i$ for d'_1 .⁴ Then we fill the second ring of exterior tiles by setting $d_2 = d'_1$ and solving the equation $p_e(d_2, d'_2) = p_i$ for d'_2 . We iterate this process until we get $d' \geq d_\infty$, where d_∞ is such that

$$\lim_{d' \rightarrow \infty} p_e(d_\infty, d') = p_i. \quad (4.13)$$

⁴ This can be formulated as finding the zero of a convex function of β'_1 with known derivative, so a modified Newton method can be applied, which ensures a solution within a given precision in both the x and y axes.

To compute this limit observe that for $d' \rightarrow \infty$ we have $\beta' \rightarrow \frac{\pi}{2}$ and then $\left(\frac{1}{\cos \beta'} - \tan \beta'\right) \rightarrow 0$, hence:

$$\lim_{d' \rightarrow \infty} p_e(d, d') = \frac{1}{\pi} \left(2d\theta + \frac{\pi}{2} - \beta - \frac{1}{\cos \beta} + \tan \beta \right) \quad \text{where} \quad \beta = \arccos \left(\frac{R \cos(d\theta)}{d} \right).$$

From the previous equation we can easily deduce that the value of d_∞ satisfying equation (4.13) is finite and satisfies

$$4 \sin(d\theta) = 2d\theta + \frac{\pi}{2} - \beta_\infty - \frac{1}{\cos \beta_\infty} + \tan \beta_\infty \quad \text{where} \quad \beta_\infty = \arccos \left(\frac{R \cos(d\theta)}{d_\infty} \right). \quad (4.14)$$

Regions in the last ring will then be unbounded, with probability $\leq p_i$. They represent parallel lines in the image plane. Figure 4.7 shows some examples of this partition of the image plane for different precision levels $d\theta$.

3.4 Final remarks

In this section we introduce some additional criteria to suppress spurious vanishing points and to eliminate the angular precision parameter $d\theta$.

Multi-precision analysis. The choice of a fixed value for the angular precision parameter $d\theta$ requires a compromise between detectability and localization error of vanishing points. We are interested in the highest possible precision level (smaller localization error in the vanishing point). On the other hand if the precision level is too fine with respect to the angular precision of the segments, the vanishing region will be hardly detected. The optimal level will approximately match the precision of the segments converging to this vanishing point, and our strategy will be to try to adjust the precision level automatically to this value. Figure 4.8 shows the value of the minimal number of concurrent lines needed for the vanishing region to be 1-meaningful, as a function of the total number of lines in the image (N) for several angular precision $d\theta$. From this figure or from simple calculations on the definition of the NFA, we find that for a total $N = 1000$ lines, we need about 300 concurrent lines to be meaningful at precision $d\theta = \frac{\pi}{16}$, whereas only 15 concurrent lines are enough at precision $d\theta = \frac{\pi}{1024}$. But we would only need 7 concurrent lines, if the total number of lines was $N = 100$. This discussion motivates the procedure described below.

As in the case of alignments, instead of fixing a single angular precision level, we will consider multiple dyadic precision levels $d\theta = 2^{-s}\pi$ for n different values of s in a certain range $[s_1, s_n]$. In our experiments $s = 4, 5, \dots, 7$ showed to be the most




Figure 4.8: Detection thresholds k for vanishing regions as a function of the total number N of lines detected in the image. For each given precision level $d\theta = 2^{-s}\pi$ (for $s = 4, 5, \dots, 10$), we construct a partition $\cup_{j=1}^N V_j = \mathbb{P}_2$ of the image plane into regions such that the probability of a line meeting each region is constant $P(G \cup V_j \neq 0 \mid G \cup \Omega \neq 0) = 4 \sin(d\theta)/\pi$. Then, given N random lines which meet the image domain Ω we compute the minimal number k of lines that should meet V_j for V_j to be ϵ -meaningful, *i.e.* for the event “at least k among N lines meet V_j ” to have an expected number of occurrences larger than $\epsilon = 1$. The figure plots k as a function of N for several precision levels s .

useful range, but this can be adjusted to the range of precision levels of the extracted segments. According to the discussion above, at each precision level $d\theta$ we should only keep those segments with a precision level no coarser than $d\theta$. Coarser segments would significantly increase N_s (thus increasing the detection threshold k) without significantly increasing the number k of lines meeting the vanishing region. Now we can apply the previously described method for all precision levels. This procedure, however, may multiply the expected number of false alarms by a factor no larger than n . So, in order to keep the false alarms rate smaller than ε we modify equation (4.6) as follows:

$$\text{NFA}(V_{j,s}) = \frac{M_s}{n} B(p_s, N_s, k). \quad (4.15)$$

The vanishing region is considered ε -meaningful if k is large enough to obtain $\text{NFA}(V_{j,s}) \leq \varepsilon$. With this definition the total expected number of false alarms from this multi-precision analysis can be easily shown to be no larger than ε . The problem is that a single vanishing point may be meaningful at several different precision levels, and we only want to keep the best explanation for it.

Local maximization of meaningfulness. When a huge number of segments meet a vanishing region $V_{j,s}$ they also meet some of the neighboring regions at the same precision level s , as well as all coarser regions $V_{j,s'} \supseteq V_{j,s}$ and some finer regions $V_{j,s''} \subseteq V_{j,s}$. Therefore, these neighboring regions too, are likely to become meaningful, but are not necessarily the best explanation. To choose among them the best explanation we introduce the following maximality concept: A vanishing region $V_{j,s}$ from a multi-precision family of partitions of the image plane is *maximal* if it is more meaningful than any other region intersecting it. More precisely, $V_{j,s}$ is maximal if:

$$\forall s' \in [s_1, s_n], \forall j' \in \{1, \dots, M_{s'}\}, \overline{V_{j',s'}} \cap \overline{V_{j,s}} \neq \emptyset \Rightarrow \text{NFA}(V_{j',s'}) \geq \text{NFA}(V_{j,s}) \quad (4.16)$$

where \overline{A} denotes the closure of a set A . Note that the condition $\overline{V_{j',s'}} \cap \overline{V_{j,s}} \neq \emptyset$ includes both neighboring regions at the same level, as well as coarser regions containing $V_{j,s}$ and finer regions contained in it.⁵

Minimum Description Length. Figure 4.9 shows all the maximal 1-meaningful vanishing regions that are detected for the photograph of a building. Clearly the first three

⁵ We used this condition instead of inclusion, because the equal probability constraint that we used to construct our partition means that regions at precision level $s+1$ cannot always be completely included in a single region at the coarsest precision level s . In this situation this “non-empty intersection”-type condition is better suited.

correspond to real orientation in the 3D scene, whereas the other three are an artificial mixture of different orientations. Observe that these mixtures are less meaningful than the original ones, because only a small portion of the segments in each direction can participate. Therefore, these artificial vanishing regions can be filtered out by an MDL criterion similar to the one we used for segments. Among all maximal meaningful vanishing regions we start a competition between them, based on the principle that each segment has to choose a single vanishing region which best explains its orientation. More precisely, a segment with supporting line l is assigned to the vanishing region $V_{j,s}$ such that $\text{NFA}(V_{j,s})$ is smallest among all regions $V_{j,s}$ met by l . Then we recompute $\text{NFA}(V_{j,s})$ for all meaningful segments using equation (4.15), with the only modification that instead of k we consider $k' \leq k$ which is the number of lines that do not only meet $V_{j,s}$, but also have been assigned to the vanishing region $V_{j,s}$. If the number of false alarms is still smaller than ε , then the vanishing region is a *maximal MDL meaningful*.

3.5 Algorithm

Summarizing, our method for vanishing region detection can be described as in algorithm 2 to obtain all maximal meaningful vanishing regions, followed by algorithm 3, if we only want the MDL maximal meaningful regions.

In order to avoid mutual exclusions, the MDL criterion may also be run iteratively as in algorithm 4. In the first iteration the V_{j_1} with lowest NFA is selected as MDL meaningful. Then k (and the corresponding NFA) is updated for the remaining meaningful V_j 's by discounting all segments meeting V_{j_1} . Thus the new NFA can only increase (see line 20) and the number of meaningful regions decreases (line 22). In the i -th iteration the i -th meaningful region V_{j_i} with lowest NFA is selected as MDL meaningful, and the remaining meaningful regions are updated by discounting from k (and the corresponding NFA) all the segments meeting V_{j_i} . The iteration stops when $\text{NFA}(V) > \varepsilon$ for all remaining regions V .

Sometimes this procedure will still miss some weak vanishing points which are “masked” by stronger vanishing points composed of much more segments. These may not be perceived at first glance, but only if we manage to unmask it by getting rid of the “clutter” in one way or another. For instance we may focus our attention into the corresponding region, or we can hide the stronger features. This unmasking mechanism can be simulated by zooming into a certain region of interest as illustrated in Figure 4.12, or by continuing our MDL iteration as follows: When no more meaningful MDL vanishing regions exist, remove all line segments that meet the already detected

Algorithm 2 Detect all maximal ε -meaningful vanishing regions**Require:** all maximal MDL 1-meaningful alignments A_1, A_2, \dots, A_N **Ensure:** all maximal ε -meaningful vanishing regions $V_{j,s}$

```

1: for all alignments  $A_i \in \{A_1, A_2, \dots, A_N\}$  do
2:   Let  $l_i$  be the supporting line for the segment  $A_i$ 
3:   Let  $d\theta_i = \sin^{-1}(\text{length}(A_i)/2)$  be the angular precision of the segment  $A_i$ .
4:   Initialize the set  $\mathcal{V}_i = \emptyset$  (vanishing regions met by  $l_i$ ).
5: end for
6: Let  $\{s_1, \dots, s_n\}$  be a range of  $n$  precision levels adapted to the angular precision of the
   alignments, i.e.
    $s_1 = \lfloor \min_i \log_2(\pi/d\theta_i) \rfloor$ ,  $s_n = \lceil \max_i \log_2(\pi/d\theta_i) \rceil$ .
7: for  $s = s_1$  to  $s_n$  do
8:   Let  $\{l_{i_1}, \dots, l_{i_{N_s}}\}$  be the supporting lines of the segments with angular precision finer
   than  $d\theta = 2^{-s}\pi$ .
9:   Compute the boundaries of the partition  $\{V_{1,s}, \dots, V_{M_s,s}\}$ .
10:  Compute the probability  $p_s$  of a random line  $G$  meeting any of the  $V_{j,s}$ 's.
11:  for  $k_0 = 0$  to  $N_s$  do
12:    Precompute the binomial tail distribution  $B(k_0, N_s, p_s) = \sum_{k=k_0}^{N_s} \binom{N_s}{k} p_s^k (1-p_s)^{N_s-k}$ .
13:  end for
14:  for  $j = 0$  to  $M_s$  do
15:    Initialize the accumulator  $k_{j,s} = 0$ .
16:    Initialize the set  $L_{j,s} = \emptyset$  (lines meeting  $V_{j,s}$ ).
17:  end for
18:  for all lines  $l_i \in \{l_{i_1}, \dots, l_{i_{N_s}}\}$  do
19:    for all vanishing regions  $V_{j,s}$  met by  $l_i$  do
20:      Increment the accumulator  $k_{j,s}$  by one.
21:      Add  $l_i$  to the set  $L_{j,s}$ .
22:      Add  $V_{j,s}$  to the set  $\mathcal{V}_i$ .
23:    end for
24:  end for
25: end for
26: for all vanishing regions  $V_{j,s}$  do
27:   Compute  $\text{NFA}_{j,s} = M_s B(k_{j,s}, N_s, p_s)$ .
28:   if  $\text{NFA}_{j,s} \leq \varepsilon/n$  then
29:     Mark  $V_{j,s}$  as meaningful.
30:   end if
31: end for
32: for all meaningful vanishing regions  $V_{j,s}$  do
33:   ismaximal = true
34:   for all meaningful vanishing regions  $V_{j',s'}$  meeting  $V_{j,s}$  do
35:     if  $\text{NFA}_{j',s'} < \text{NFA}_{j,s}$  then
36:       ismaximal = false
37:       break
38:     end if
39:   end for
40:   if ismaximal then
41:     Mark  $V_{j,s}$  as maximal meaningful.
42:   end if
43: end for

```

Algorithm 3 Detect all MDL maximal ε -meaningful vanishing regions

Require: all maximal ε -meaningful vanishing regions $V_{j,s}$,

and the corresponding meeting lines $L_{j,s}$.

Ensure: all maximal MDL ε -meaningful vanishing regions $V_{j,s}$

- 1: Initialize $k'_{j,s} = 0$ for all input $V_{j,s}$'s.
 - 2: **for all** lines l_i **do**
 - 3: Compute $(j_i, s_i) = \arg \min\{\text{NFA}_{j,s} : V_{j,s} \text{ is maximal meaningful and } l_i \in L_{j,s}\}$
 (this is the most meaningful vanishing region meeting line l_i).
 - 4: Increment k'_{j_i, s_i} by one
 - 5: **end for**
 - 6: **for all** maximal meaningful vanishing regions $V_{j,s}$ **do**
 - 7: Compute $\text{NFA}'_{j,s} = M_s B(k'_{j,s}, N_s, p_s)$
 - 8: **if** $\text{NFA}_{j,s} \leq \varepsilon/n$ **then**
 - 9: Mark $V_{j,s}$ as MDL maximal meaningful.
 - 10: **end if**
 - 11: **end for**
-

vanishing points V_{j_1}, \dots, V_{j_i} . Thus the total number N of segments will decrease and so will do $\text{NFA}(V_j) = MB(p, N, k)$ in equation (4.6) or line 27 of algorithm 2. Thus some vanishing points may become meaningful again and we can restart the previous iteration. This iteration allows to distinguish a first group of features that are meaningful on its own right, from a second group which is only detected in the absence of the first group's masking effect (see algorithm 5).

The complexity of algorithm 2 for N lines and M tested vanishing regions is dominated by the line-cell intersections in line 19, which requires $o(N\sqrt{M})$. This is due to the fact that all M regions are delimited by only $o(\sqrt{M})$ lines and circles. Thus we only need to compute these $o(\sqrt{M})$ line-line or line-circle intersections for each of the N lines. Then, from the intersection points we can immediately retrieve the vanishing regions that are concerned. The memory requirements are however quite demanding because we need hash tables or a similar data structure in order to efficiently handle each of the M sets $L_{j,s} \in 2^{\{l_1, \dots, l_N\}}$, and for each of the N sets $\mathcal{V}_i \in 2^{\{V_1, \dots, V_M\}}$.

The simple MDL algorithm 3 takes $o(M'' + NM')$ operations where $M'' \ll M$ is the total number of maximal meaningful regions and $M' \approx \sqrt{M}$ is the average number of regions met by a single line. However, this number could be brought down to $o(M'' + N)$ if we use a sorted data structure for \mathcal{V}_i , which does not significantly penalize

Algorithm 4 Detect all MDL maximal ε -meaningful vanishing regions avoiding mutual exclusions

Require: \mathcal{M} = the set of all maximal ε -meaningful vanishing regions $V_{j,s}$.

The corresponding set of meeting lines $L_{j,s}$, number of false alarms $\text{NFA}_{j,s}$ for all $V_{j,s} \in \mathcal{M}$.

The set \mathcal{V}_i of vanishing regions met by l_i , for all $i = 1, \dots, N$.

Ensure: The set \mathcal{M}' of all MDL maximal ε -meaningful vanishing regions $V_{j,s}$

```

1: for all maximal meaningful vanishing regions  $V_{j,s} \in \mathcal{M}$  do
2:   Initialize  $\text{NFA}'_{j,s} = \text{NFA}_{j,s}$  and  $L'_{j,s} = L_{j,s}$ .
3: end for
4: for all lines  $l_i$  do
5:   Initialize  $\mathcal{V}'_i = \mathcal{V}_i \cap \mathcal{M}$  (we only care about maximal meaningful vanishing regions).
6: end for
7: Initialize  $\mathcal{M}' = \emptyset$ 
8: while  $\mathcal{M} \neq \emptyset$  do
9:    $(j, s) = \arg \min \{ \text{NFA}'_{j,s} : V_{j,s} \in \mathcal{M} \}$  (this is the most meaningful vanishing region in
     $\mathcal{M}$ ).
10:  Add  $V_{j,s}$  to  $\mathcal{M}'$ .
11:  for all lines  $l_i \in L_{j,s}$  do
12:    for all  $V_{j',s'} \in \mathcal{V}'_i \cap \mathcal{M}$  do
13:      Remove  $l_i$  from  $L'_{j',s'}$  (set of lines meeting  $V_{j',s'}$ ).
14:      Remove  $V_{j',s'}$  from  $\mathcal{V}'_i$  (set of maximal meaningful vanishing regions met by  $l_i$ ).
15:    end for
16:  end for
17:  Remove  $V_{j,s}$  from  $\mathcal{M}$ .
18:  for all  $V_{j',s'} \in \mathcal{M}$  do
19:    Update  $k'_{j',s'} = \#L'_{j',s'}$  (it may have decreased).
20:    Update  $\text{NFA}'_{j',s'} = M_s B(k'_{j',s'}, N_{s'}, p_{s'})$  (it may have increased).
21:    if  $\text{NFA}'_{j',s'} > \varepsilon/n$  then
22:      Remove  $V_{j',s'}$  from  $\mathcal{M}$  (it is no longer meaningful).
23:    end if
24:  end for
25: end while

```

Algorithm 5 Detect all unmasked MDL maximal ε -meaningful vanishing regions

Require: The set of detected lines $\mathcal{L} = \{l_1, \dots, l_N\}$.

Ensure: A series of sets $\mathcal{M}'_0, \dots, \mathcal{M}'_m$, such that

- \mathcal{M}'_0 contains the set all MDL maximal meaningful vanishing regions, and for $h \in [1, m]$
- \mathcal{M}'_h contains all vanishing regions that become MDL maximal meaningful, when we remove all lines contributing to a vanishing region in all previous levels $0, \dots, h - 1$.

- 1: Initialize $m = 0$.
 - 2: Initialize $\mathcal{M}_m = \text{Algorithm 2}(\mathcal{L})$.
 - 3: **while** $\mathcal{M}_m \neq \emptyset$ **do**
 - 4: $\mathcal{M}'_m = \text{Algorithm 4}(\mathcal{M}_m)$.
 - 5: Remove $\bigcup_{V_{j,s} \in \mathcal{M}'_m} L'_{j,s}$ from \mathcal{L} .
 - 6: Increment m by one.
 - 7: $\mathcal{M}_m = \text{Algorithm 2}(\mathcal{L})$.
 - 8: **end while**
-

the complexity of insertions.

The more elaborate MDL algorithm 4 which takes about $o(NM'')$ per iteration where $M'' \ll M$ is the number of meaningful regions that have been detected at that iteration.

The number M is fixed and equal to 270360, 67828, 17195, when the finest angular precision level is respectively $\frac{\pi}{1024}$, $\frac{\pi}{512}$, $\frac{\pi}{256}$, *i.e.* we have $M \approx p^{-2}/4$ where p is the precision level (inverse of the number of distinguishable orientations).

An optimized version of this software runs on a 1GHz Pentium processor in 0.16 seconds for $N = 64$, 0.24s for $N = 128$, 0.41s for $N = 256$, 0.79s for $N = 512$, up to 1.88s for $N = 1024$. In an image of size 512×512 we usually detect a few hundred alignments, so the running time of the vanishing point detector part is negligible with respect to the alignment detection module which does the bulk of the work. The complexity of the latter module is $o(N^4)$ for an image of size N , and the running time for $N = 512$ is about 40 seconds, also on a 1GHz Pentium processor.

4 Experiments and Discussion

Figures 4.9, 4.11 and 4.13 show the results of applying our algorithm for vanishing point detection on several images.⁶ In most cases consisting of man-made environments the most relevant orientations are detected, without any false alarms. Figure 4.10 illustrates the need of the MDL criterion in order to filter out artificial vanishing points that may appear when the real vanishing points are extremely meaningful. Note that after MDL (figures 4.9(b) to 4.9(d)) we only get the main three directions (two horizontal and one vertical).

Figure 4.12 illustrates the masking phenomenon. Here the less meaningful directions corresponding to the wall are “masked” by the many segments in the horizontal and vertical directions, but it can be “unmasked”. See the figure captions for a more detailed explanation.

Finally Figure 4.13 shows the limitations of the proposed method when applied to natural images not containing vanishing points (see caption for details). This and other similar experiments further enforce the conclusion in [64] on the importance of addressing the conflicts between Gestalts. Indeed if we were able to resolve the conflict between the alignment and the curved boundaries Gestalts we would eliminate many “false” line segments and thus further reduce the number of false alarms in the vanishing point detection phase. Our experiments suggest that this approach might be complementary to (and in certain cases better adapted than) other approaches to reducing spurious responses rather based on joint Gestalts at the vanishing point detection level.

It is quite difficult to build an experimental setup which allows to fairly compare our method with previously proposed ones. The reason is that our assumptions are quite different here, since we do not treat the same problem: whereas most previous works [120, 146, 156, 160] look for joint Gestalts that combine 3D parallelism with some other property, whereas here he try to push the pure partial Gestalt of 3D parallelism to its limits.

An exception is the recent work in [23], which only relies on 3D parallelism and has been shown to produce highly accurate vanishing points, but assumes knowledge of

⁶ In all our experiments we used $\varepsilon = 1$, for coherence with previous works. But we could also have used a much smaller value, since in all the examples presented here all real vanishing points have $NFA < 0.0001$. Furthermore $\varepsilon = 1$ means that we can expect in average one false vanishing point in a random image, which is quite high with respect to the reduced number of vanishing points we usually find in real images.

the camera calibration parameters and omnidirectional images, which is not exploited by our method. The importance of this knowledge is not thoroughly discussed in [23], but it was crucial in [120] in order to reduce spurious responses. The work in [23] relies on a Hough transform as in [120] in order to determine the number of vanishing points and is therefore prone to the same sensitivity to the internal calibration parameters. For this reason the method proposed in [23] can be considered as complementary to our method. In fact our method could be used either in the initialization its step to determine the number and approximate positions of vanishing points more reliably, or as a validation step to reduce the number of false alarms.

5 Conclusions and Perspectives

We presented a fully automatic vanishing point detector which requires no a priori information, and discussed its applications to the calibration of internal and external camera parameters, single and multiple view photogrammetry, and registration, among others. In this section we discuss some ideas for further improvement or generalization of the proposed method as well as its relationships work with other areas of research in computer vision, that we consider worth exploring.

Gestalt collaborations, conflicts and masking. The experiments commented in the previous section showed how different variants of the MDL criterion can be used to deal with common problems in feature detection, such as multiple or spurious responses, as well as the masking phenomenon. We also suggested how it could be used to reduce the number of false alarms of the proposed vanishing point detector when applied to images that actually contain no vanishing points.

More generally, the MDL criterion can be viewed as a technique for solving conflicts between competing Gestalts, and other similar techniques can be applied as well. The importance of addressing the collaborations and conflicts of grouping laws, and the masking issue has been remarked long ago within the Gestalt community [108, 173], and more recently in the computer vision community [63].

The two methods presented here actually only address the collaboration between objects in the same category (either edgels in the case of alignments, or segments in the case of vanishing points). The experiments shown here further support the conclusion stated in [63] that in order to come closer to the full programme of Wertheimer and Kanisza, we need to move up one more level in the hierarchy and also consider collab-

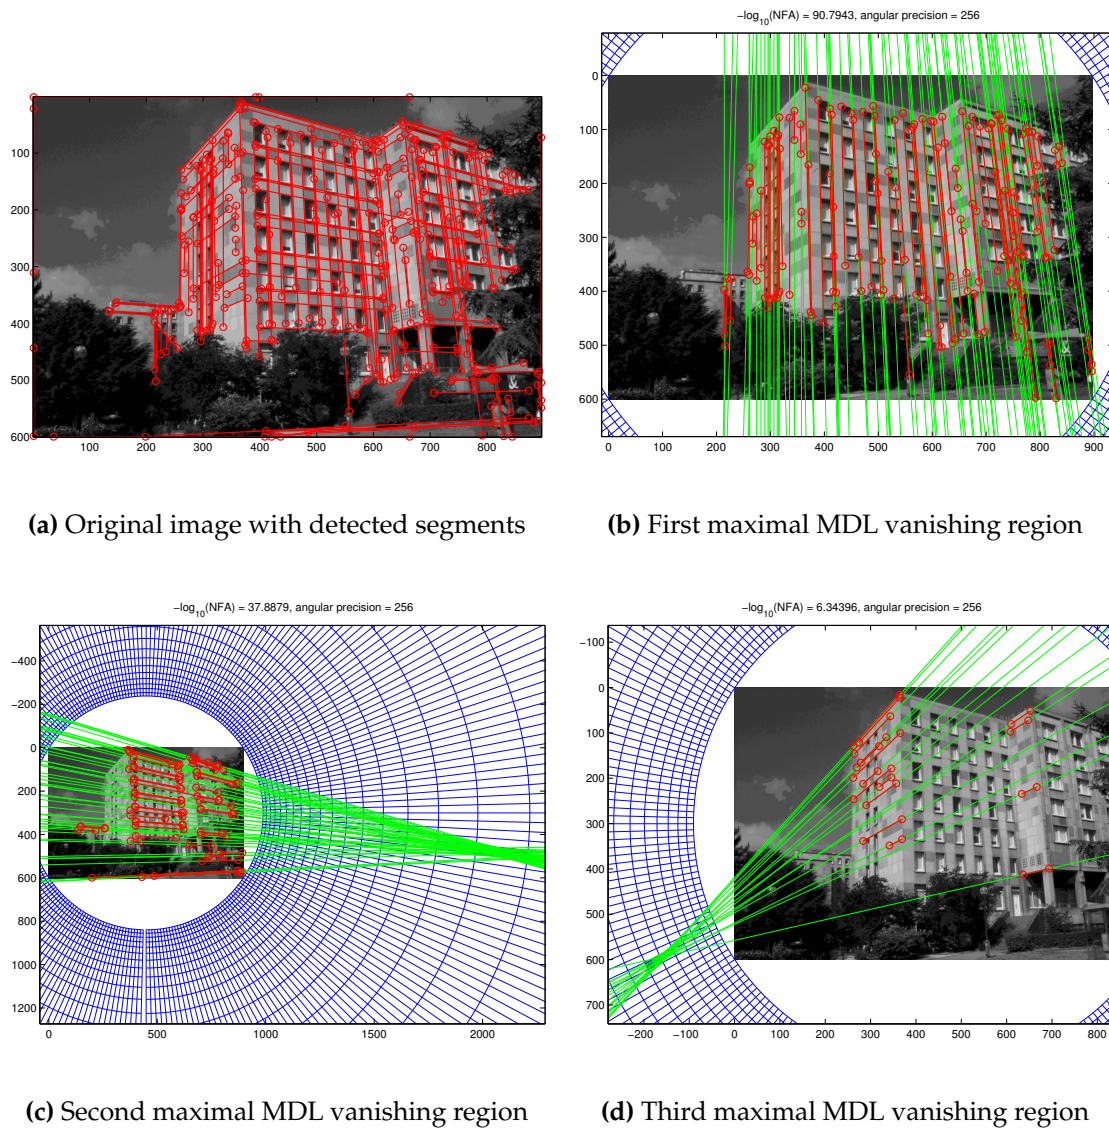


Figure 4.9: Detected line segments for a building image and the only three maximal MDL vanishing points that are detected. They correspond to the two horizontal orientations and to one vertical orientation. Note that no orthogonality hypothesis was used, thus it can be used a posteriori in order to calibrate some camera parameters. For each vanishing point we only display the segments that contributed to this point at the automatically chosen precision.

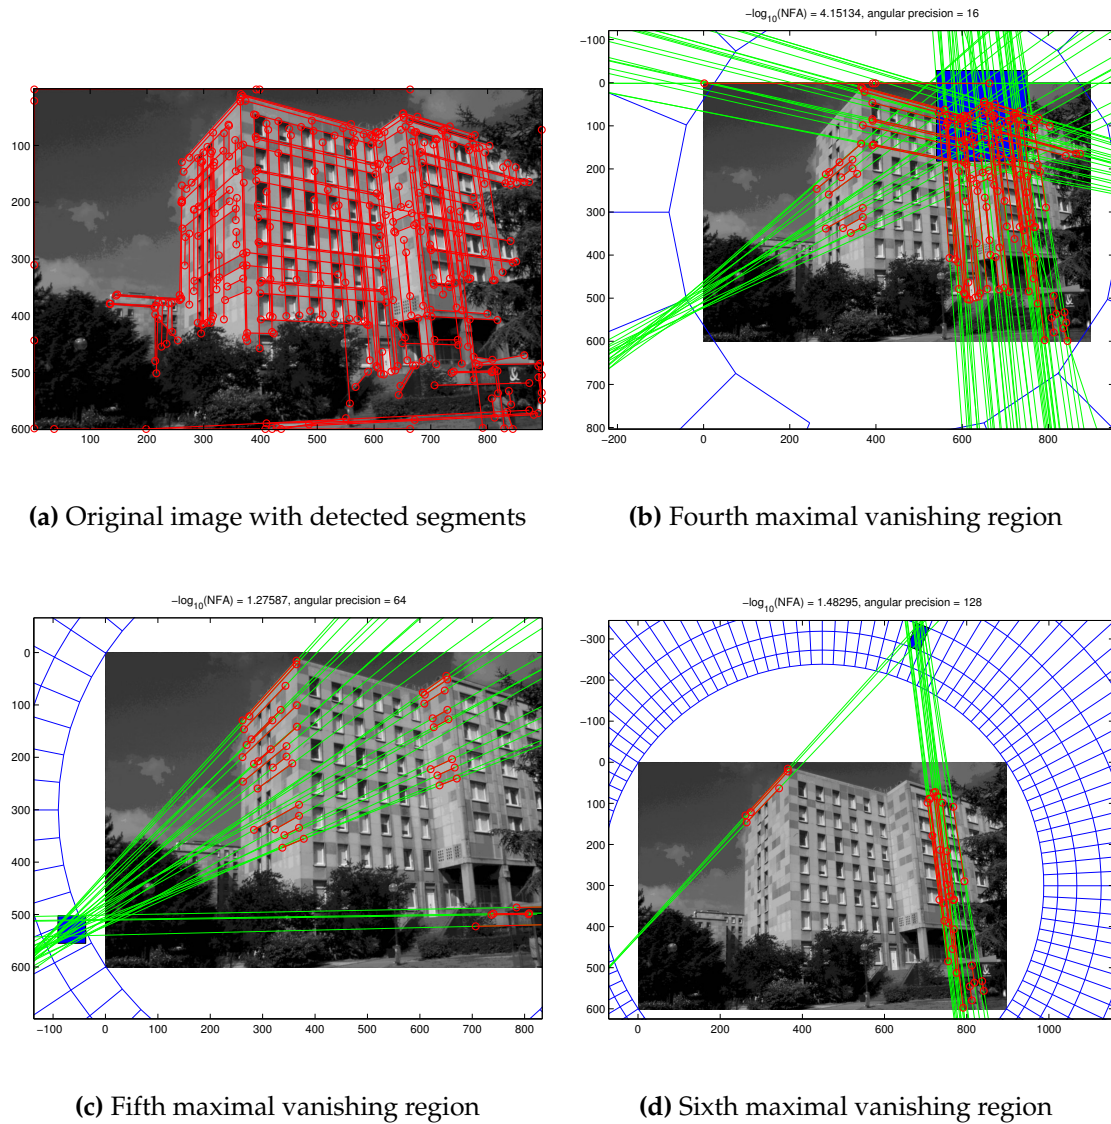


Figure 4.10: Before applying the MDL criterion some spurious vanishing regions remain. Note that they arise from mixtures of real vanishing regions, and that they are significantly less meaningful and less precise than the real vanishing regions. Therefore, during MDL, most segments vote for the real vanishing region instead of these “mixed” ones, so that after MDL their number of false alarms decreases and they are no longer meaningful.

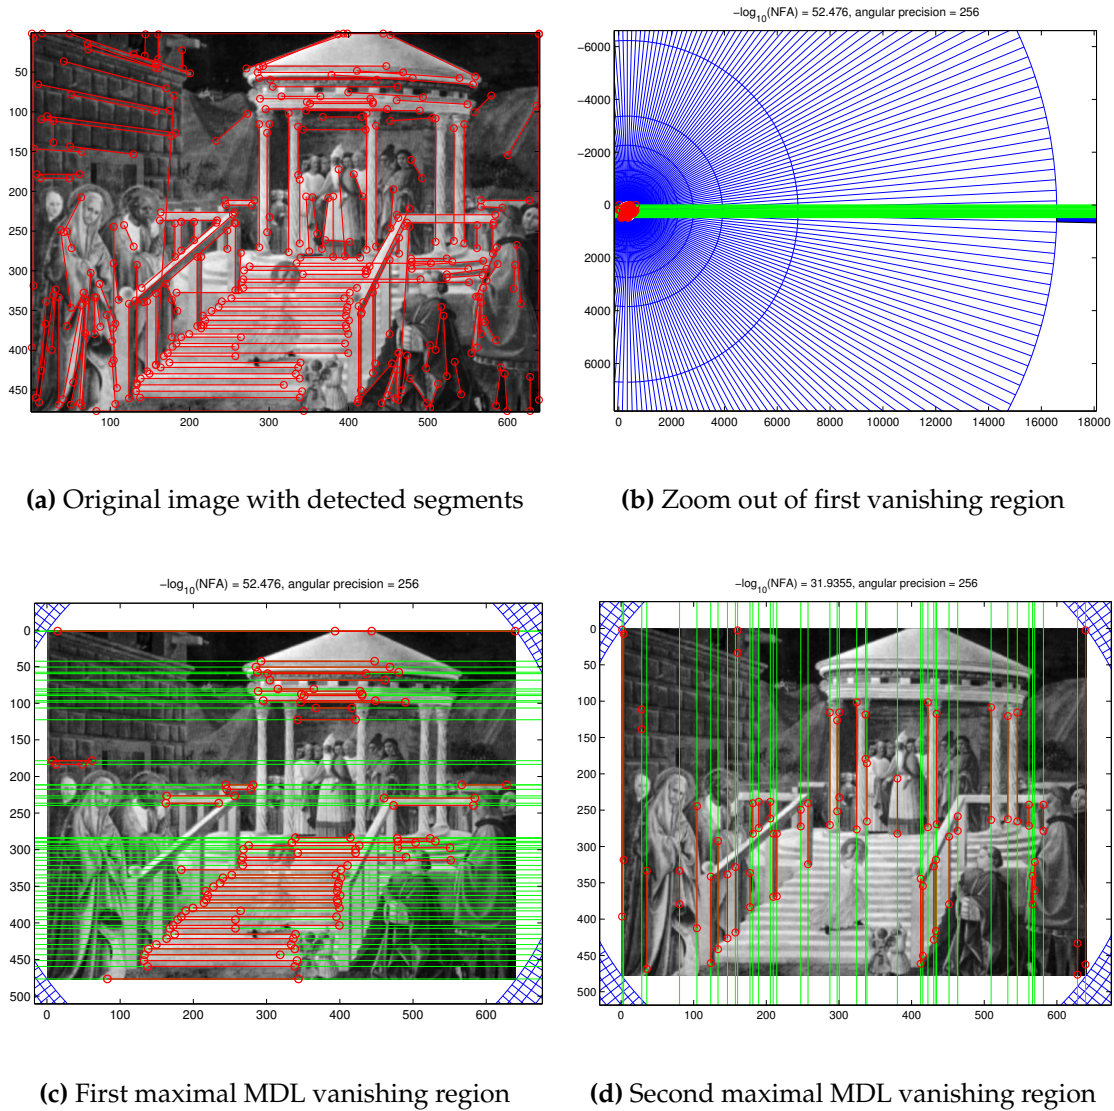
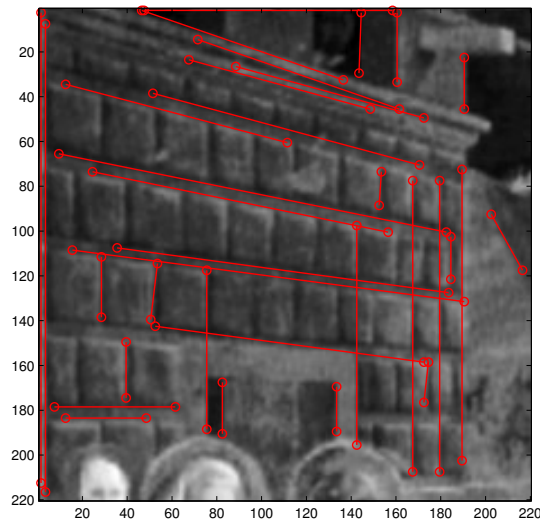
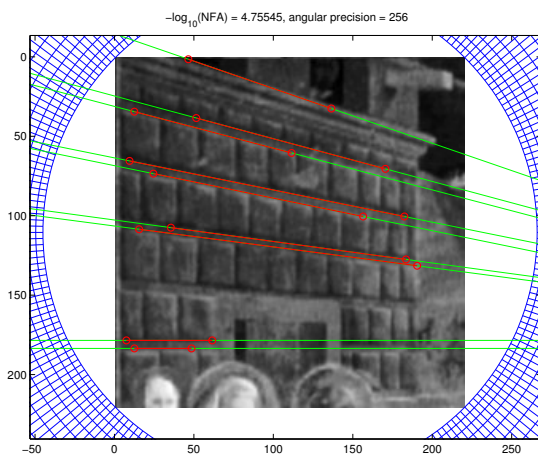


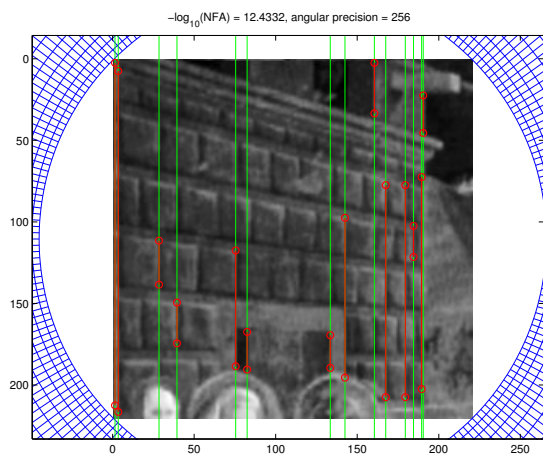
Figure 4.11: Detected line segments for an image of a painting by Uccello and the only two maximal MDL vanishing points that are detected. Note that the vanishing points corresponding to the oblique wall and the staircase are missed. This is due to the fact that both the alignment detection and the vanishing point detection are global, and the less meaningful segments and vanishing points are masked by the more meaningful horizontal and vertical orientations (see figure 4.12)



(a) Original image with detected segments




(b) First maximal MDL vanishing region

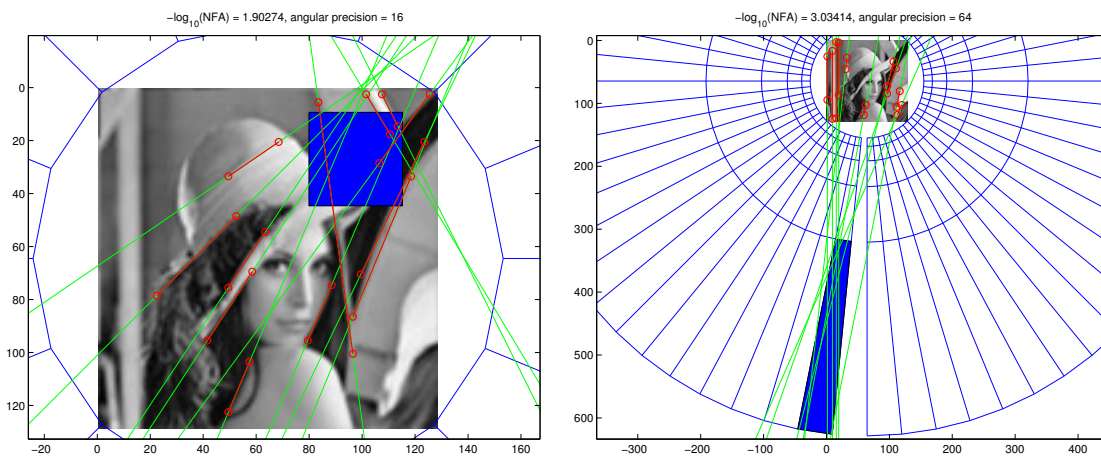


(c) Second maximal MDL vanishing region

Figure 4.12: *Illustration of the “masking” phenomenon.* When we select the wall subimage in the previous figure, more alignments are detected, and a new vanishing point that was masked in the global image becomes meaningful. This is due to two cooperating effects. First the masking phenomenon at the alignment detection level means that we detect in this subimage more meaningful segments than in the global image. Secondly, at the vanishing points detection level, the total number of segments is smaller, which means that the minimal number of concurrent lines for a vanishing region to become meaningful k is also small. A similar result may be obtained by restarting the MDL iteration a second time with the remaining segments after all MDL meaningful vanishing points have been detected and the contributing segments removed.



(a) Original image with detected segments



(b) First maximal MDL vanishing region

(c) Second maximal MDL vanishing region

Figure 4.13: *Accidental vanishing points.* When applied to images of man-made environments which actually contain vanishing points, the method very rarely detects accidental vanishing points. But this does happen in natural images in which we do not perceive such vanishing points. Here we show one of the worst such examples that we found in our experiments. In this case the detected vanishing points are probably not perceived because they are made up mostly of segments that are not perceived as straight lines in the first place. Many of these segments would be better explained as meaningful curved boundaries, and therefore will never give rise to vanishing points. Hence the false alarms in the vanishing points detection phase are here to some extent the result of some special kind of false alarms in the alignment detection phase. Further experiments on natural images showed this kind of false alarms of vanishing points (due to some false alarms in line segments which are actually curved boundaries) to be the most prominent one.

orations and conflicts between *different* grouping laws (such as alignments and curved boundaries in the case of the Lena image). They also show possible mechanisms based on the MDL for dealing with conflicts between grouping laws and the masking phenomenon. These mechanisms could eventually be generalized to other partial or joint Gestalts.

The question of efficiency and false negatives. The feature detection methods presented in this chapter are focused on minimizing *false positives* (*i.e.* false alarms). They can be viewed as *verification* methods, that given a feature reliably decide whether the feature is significant (and thus perceived) or not. In order to obtain a low false alarms rate with no missing feature we just sample the whole feature space, and verify every single feature. We could say that these methods concentrate on verification, rather than search.

In this sense there is a nice complementarity with the work of Amit and Geman [20] which on the contrary focuses on the *search* or visual selection process, rather than the verification step. This search is guided by a computational economy principle, which turns up to be tightly related to minimizing *false negatives* (*i.e.* misdetections). More specifically he builds search strategies, often in the form of a coarse-to-fine tree, in which he can prune with inexpensive computations large branches of the tree having a negligible probability of containing the searched feature.

A combination of Geman's visual selection methods with the a-contrario models used here and in [60, 62, 63, 64], would possibly lead to very efficient methods with both a low false positive and a low false negative rate.

Without pretending to develop an optimized search strategy as in Geman's papers, we present in the following paragraphs a few ideas that could be used to make the proposed methods more efficient.

Improving efficiency of the alignment detector. In order to improve the efficiency of the alignment detector which is currently the bottleneck of our vanishing point detector we can think of different ways to safely and quickly prune the search tree, *i.e.* determining with little computations a subset of segments, such that the complement has zero or almost zero probability of containing a meaningful segment.

Gradient thresholding and Canny points. Instead of considering all pairs of points as alignment candidates we can search among pairs of points with gradient above the threshold in equation (4.4), or even Canny points.

$N \setminus n$	2	4	8	16	32	64
64	24	12	8	6	5	4
128	28	14	10	7	6	5
256	32	16	11	8	7	6
512	36	18	12	9	8	6
1024	40	20	14	10	8	7

Table 4.1: Minimal number l of independent points for a segment to be meaningful at angular precision equal to π/n , in an image of size $N \times N$. The actual length of the segment is $2l$ pixels.

Hough Transform. Pushing the previous idea a bit further, we can first select the set of candidate supporting lines, from a Hough transform of gradients above the threshold (we can also use the gradient orientation and its angular precision $\frac{q}{|Du|}$ to limit the number of bins to which the point contributes). Now we can compute the minimal number of points needed in a bin, for that line to contain at least one meaningful segment. Observe that at precision p , a segment of length at least $l = \min\{l : MB(p, l, l) < \varepsilon\}$ is needed to define a meaningful alignment. For the common precision of $p = 1/16$, this minimal length varies linearly between 12 and 20 (when the size of the image varies exponentially between 128 and 1024), but it doubles when the precision is $p = 1/4$. However points with weak gradient (according to equation (4.4)) don't count. So a simple strategy consists of searching for significant segments only along lines containing at least l aligned points with gradient above the threshold. Such lines can be easily extracted by means of a Hough transform.

MDL. The MDL criterion itself may be used, in conjunction with the Hough transform in the previous paragraph, in order to further prune the search tree in a safe way, if we only look for MDL maximal segments. Assume a particularly long and significant segment A has been detected. Then we can proceed the search in two different ways:

Less meaningful segments B cannot take away an aligned point from A . Thus the number of points contributing to the Hough transform decreases by the same amount as aligned points in A , and so does the number of candidates supporting lines, reducing the search space.

More meaningful segments C (i.e. with $NFA(C) \leq NFA(A) \leq \varepsilon$) will require even

more aligned points, so the Hough transform thresholds in table 4.1 will increase (since we are looking now for $NFA(\varepsilon)$ -meaningful segments), thus reducing the search space too.

In conclusion each newly detected segment allows us to reduce the search space in two different ways. The larger and meaningful the segment is, the more it will reduce this space. A reasonable way to organize the search using these two principles may be as follows:

Start with the line containing the largest number of points satisfying the gradient threshold equation (4.4) (this is just the global maximum of the Hough transform). Search the most meaningful segment A along this line. This one has chances to be also very long.

Update the Hough transform thresholds for $NFA(A)$ -meaningfulness, and look for any more meaningful segment A' . A first good guess may be the second global maximum of the Hough transform (if it is still $NFA(A)$ -meaningful).

We iterate this procedure until we get the most meaningful segment C . At this point C is certainly MDL maximal meaningful, and it certainly owns all the aligned points in C . So we may definitely discard all aligned points in C from further processing, *i.e.* now we look for less meaningful segments B , so we can update the Hough transform accordingly, and we may restart the iteration. Of course, a non-trivial data structure is needed to avoid repeating the same computations more than once.

Improving precision of detected segments. We can use at least two mechanisms to further increase the angular precision of the detected line segments.

Local search. Use an appropriate image interpolation based on the available knowledge of the image formation, to obtain a continuous version of Du . The interpolation method should also allow fast computation of the gradient. A good compromise might be achieved by bilinear interpolation of an oversampled image. Then, instead of considering discrete segments (where the pixel centers are not exactly aligned), consider the alignment of regularly spaced points along the continuous line segment. Finally, once an MDL maximal meaningful alignment A has been found, perform a gradient descent of $NFA A'$ among all segments A' whose endpoints lie within a neighborhood of less than one pixel from A . One or two iterations of gradient descent should be enough to obtain subpixel accuracy in the location of the segment's endpoints.

Grouping of aligned segments. The current procedure, due to its global nature, may sometimes “fill in the gaps” between two aligned segments. Even if the points in the gap are not aligned, if the gap is small enough, the combined segment may be more meaningful than each of the two segments alone. However, if the gap is too big, the aligned segments will not be grouped into a single segment. But we want the detected segments to be as long as possible in order to increase their angular precision given by $\sin^{-1} \frac{1}{l}$, where l is the length of the segment. Hence we could add a third grouping law allowing the detected segments to group together if they are significantly aligned with aligned orientations.

Improving detectability of segments. Instead of MDL with dilated segments, we can consider segments of different widths and detect the most meaningful among all widths. This would not only solve the multiple response problem as well, but it may also detect low-contrast blurred edges as well, which currently escape our detector (note that for wide segments, the number of points increases not only with the length, but also with the width of the segment). This way, the angular precision $\sin^{-1} \frac{w}{l}$ can be computed automatically as well, but it depends now also on the detected width w of the segment. A potential problem with this approach is that oriented textured regions may appear as a large almost square block, if we use gradient orientations modulo π . This problem may be partially solved by considering “signed” orientations, *i.e.* modulo 2π . This way the two sides of a ridge are distinguished as distinct edges.

Link to Sampling and Information theory. The number of false alarms defined in equation (4.3) for the alignment detector involves the length l of the line segment and the number k of aligned points. But l and k are not measured in terms of pixels; instead we count only one every two pixels on a line. The argument that led to this choice in [61], is that we need at least a 2×2 neighborhood to compute a gradient orientation, if we use the usual $[-1, 1]$ stencil for finite differences. Therefore when we compute gradient orientations on an image where the pixels are all independent we obtain independent measures only at the pixels with *e.g.* even coordinates, *i.e.* one every two points on a line, or one every four points in the image. In practice this is implemented as follows. Assume we are computing the NFA of a segment A of length L pixels, and we count K aligned points. Then we define $l = L/2$ and $k = K/2$, and as usual $\text{NFA}(A) = N^4 B(p, l, k)$.

Now we could argue that if the image is band-limited, we could, by Fourier interpolation, obtain uncorrelated measurements of gradient orientation at every single pixel

(though not necessarily independent). This is however not true for real images, because it is impossible to construct sampling instruments which produce band-limited, non-aliased images where all pixels are independent, as it has been suggested in chapter 1. In fact if we want to have images where aliasing effects are reasonably controlled, with a standard CCD array the ratio $R_{\text{eff}}/R_{\text{nom}}$ which can be roughly interpreted as the mean distance (measured in pixels, *i.e.* as a factor of the sampling period) between two independent points is about 1.2. With the more efficient Supermode instrument this ratio is about 1.1. More interestingly when images are acquired by an optical system with a wide field of view there is often a radial geometric distortion. After the image has been rectified to avoid straight-lines to become curved, resolution is finer on the center than near the corners, and the $R_{\text{eff}}/R_{\text{nom}}$ ratio changes from one point to another. In any case, if the characteristics of the acquisition system are known, then the alignment detector could be adapted to them by defining $l = L/C$ and $k = K/C$, where $C = C'R_{\text{eff}}/R_{\text{nom}}$ represents the distance between two independent gradient orientations and $C' \in [1, 2]$ is a constant that depends on the interpolation method used to estimate the gradient orientation from the image pixels. This corresponds to changing our a-contrario model (from which we seek large deviations) from assuming that the *sampled image* is white noise, to assuming that the *observed landscape, before sampling*, is a white noise, and taking into account the degradations introduced by the acquisition system.

If such informations about the instrument are not taken into account and corrected, then we can still use the original a-contrario model, but segments may become less meaningful, not detected or detected at a precision different from the optimal one, or with less angular precision. In this situation it is more advisable to use a large range of angular precisions during alignment detection. As an example, in an image taken in the conditions previously described (a rectified image from a wide field of view camera, sampled by a CCD array), not enough segments were detected at the usual precision $p = 1/16$. Expanding the range of precision levels from $p = 1/4$ to $p = 1/16$ we obtained many more segments and both vanishing points were detected.

Even when the exact acquisition conditions are known, a trade-off must be solved concerning to what an extent aliasing and noise should be removed, and the transfer function deconvolved from the raw image before alignment detection. Lower aliasing and noise thresholds mean that we can afford higher angular precision levels, thus contributing to higher detectability. On the other hand lower thresholds also mean

a smaller optimal reciprocal cell and hence a coarser effective resolution, so we have less independent points, and detectability decreases. Furthermore, restoration methods may themselves introduce certain correlations which may lead at the end to false detections. The trade-off does not only concern detectability but also efficiency, for two reasons. First if the effective resolution is known then the image may be downsampled without loss of information by a factor z , thus reducing the running time of the algorithm by a factor z^4 . Secondly, if we master the kinds of errors in our image, then we can also reduce the ranges of precision levels that should be tested, thus reducing time of computation as well.

Improving the efficiency of the vanishing point detector. In certain situations where the number of detected segments is very small it may not be worth building a complete tiling of the plane into vanishing regions. It may not be worth both from the computational and from the detectability point of view, as we justify below.

In such cases it is more interesting to consider the vanishing region family $\{V_{i,j}\}_{i<j}$ composed of the intersections

$$V_{ij} = \text{Cone}(A_i) \cap \text{Cone}(A_j)$$

of the uncertainty cones for all pairs of segments.⁷ Here we use the same principle as has been proposed in [63] to detect strips containing a large number of points, or in the previous paragraph to improve the efficiency of the alignment detector. The total number of vanishing regions is $M' \in [N(N-1), 4N(N-1)]$ roughly proportional (up to a constant smaller than 4) to the square of the total number N of detected segments. But all M cells are delimited by a total of $2N \approx 2\sqrt{M'}$ lines. So the line-cell intersections can be computed as in the previous case in $O(N\sqrt{M'})$ operations. So the method is useful when $M' < M$, *i.e.* when the number of segments $N < \sqrt{M}$ is small with respect to (the square root of) the number of cells. But since the number of cells M in the partition method is roughly proportional to n^2 where n is related to the maximal precision level through $d\theta = \frac{\pi}{n}$, we find that this method is useful from the computational point of view whenever the number of segments $N < n$ is smaller than the number of orientations that determines the maximal precision. Also when $N < n$ we have $M' < M$ and the number of tests is reduced, so the number of false alarms $\text{NFA}(V) = MB(p, N, k)$ decreases, and the vanishing regions become also more detectable.

⁷Occasionally these vanishing regions may be composed of two distinct connected components. In such a case the computation of the probability of a line meeting such a region is not a solved problem, and we may be better off splitting this region into two.

Now both N and n are usually linearly related to the size of the image and we often have $N \approx n$, but the variations may be very large between images of natural environments and of man-made environments, so a sensible choice between both methods cannot be fixed a priori. Thus an adaptive algorithm which selects the cheapest of both methods based on the relationship between N (the total number of detected segments) and n (the maximal precision of a detected segment), would be needed to accelerate the performance of vanishing point detection in extreme cases.

Appendix

A Proof of the main results on integral geometry.

The first case $K_1 \subseteq K_2$, in proposition 11 is a direct consequence of the following lemma

Lemma 11. *Given a convex set K , the measure of a random line G meeting the set equals its perimeter.*

Proof. Since our measure has been chosen to be translation invariant we may assume without loss of generality that the origin of our coordinate system O is inside K . In such a case, the convex set K (containing the origin of the coordinate system) is completely described by the support function, see figure 4.3(b))

$$p(\theta) := \sup\{\rho : G(\rho, \theta) \cap K \neq \emptyset\}. \quad (4.17)$$

Then we observe that the arc-length parameter s of ∂K can be written after some calculus as (see [155, section 1.1.2, pages 2–3])

$$ds = (p(\theta) + p''(\theta)) d\theta, \quad (4.18)$$

which allows to compute the length L of ∂K as:

$$L = \int_{\partial K} ds = \int_0^{2\pi} p(\theta) d\theta \quad (4.19)$$

because the integral of $p''(\theta)$ is zero.

But by definition of the support function, the set of lines meeting K with orientation θ are described by $\rho \in [0, p(\theta)]$. So the measure of all lines intersecting K , is also:

$$\mu[G \cap K \neq \emptyset] = \int_{G \cap K \neq \emptyset} dG = \int_0^{2\pi} \int_0^{p(\theta)} d\rho d\theta = \int_0^{2\pi} p(\theta) d\theta \quad (4.20)$$

This shows that the measure of all lines intersecting K is the length L of ∂K , i.e. $\mu[G \cap K \neq \emptyset] = L$. \square

The reader can easily check that the first case in proposition 11, follows directly from the lemma for $K = K_2$, since the inclusion $K_1 \subseteq K_2$ means that the event $G \cap K_1 \neq \emptyset$ is implied by the event $G \cap K_2 \neq \emptyset$. To show the second case a generalization of this lemma is required. Its proof may be found in [155, section 1.3.2, page 31] .

Lemma 12. *Let C be a piecewise differentiable curve in the plane that has a finite length L . Let G be a random line in the plane, which intersects the curve C at n different points. Then*

$$\int n \, dG = 2L. \quad (4.21)$$

Observe that lemma 11 is a particular case of this lemma when $C = \partial K$, since in that case

$$n = \begin{cases} 2 & \text{if } G \cap \overset{\circ}{K} \neq \emptyset \\ 1 & \text{if } G \cap \overset{\circ}{K} = \emptyset \text{ and } G \cap K = \emptyset \\ 0 & \text{if } G \cap K = \emptyset. \end{cases}$$

Put another way $n = 1$ in the limit case when G is tangent to K but this case has clearly zero measure, and therefore

$$\int n \, dG = \int_{G \cap K \neq \emptyset} 2 \, dG = 2\mu[G \cap K \neq \emptyset].$$

Now we consider case 2 in proposition 11, namely the case when $K_1 \cap K_2 = \emptyset$. Following [155, section 1.3.3, pages 32–34] we consider four curves, namely the two perimeters $\partial K_1, \partial K_2$, and the interior and exterior envelopes C_e and C_i . Let \mathcal{C} be the concatenation of these four curves.

It turns out that the three different cases where G meets both K_1 and K_2 , only one of them, or neither of them but still separating them can be characterized by the number n of intersections of G with the curve \mathcal{C} :

$$n = 10 \Leftrightarrow G \text{ meets both } K_1 \text{ and } K_2$$

$$n = 6 \Leftrightarrow G \text{ meets either } K_1 \text{ or } K_2 \text{ but not both.}$$

$$n = 4 \Leftrightarrow G \text{ separates } K_1 \text{ and } K_2 \text{ without meeting either of them.}$$

In the last equation, to separate, means that K_1 and K_2 are on different half-planes with respect to the line G . Let's call m_{10}, m_6 and m_4 the measures of the sets of lines in each case, and

$$m_6 = m'_6 + m''_6 \quad (4.22)$$

to make the difference between the lines that meet only K_1 from those meeting only K_2 . Then from simple disjoint union arguments (up to a set of zero measure) we can establish the following equations

$$m'_6 = L_1 - m_{10} \quad (4.23)$$

$$m''_6 = L_2 - m_{10} \quad (4.24)$$

$$m_4 + m_6 + m_{10} = L_e \quad (4.25)$$

The last equation follows from the application of lemma 11 to the convex set delimited by C_e , which is met by G in all three cases. In addition, if we apply lemma 12 to the concatenated curve \mathcal{C} we get

$$4m_4 + 6m_6 + 10m_{10} = 2(L_1 + L_2 + L_i + L_e) \quad (4.26)$$

Finally equations (4.22) through (4.26) define a linear system of 5 equations in the 5 unknowns $m_4, m_6, m'_6, m''_6, m_{10}$. Solving for m_{10} we get

$$\mu[G \cap K_1 \neq \emptyset \text{ and } G \cap K_2 \neq \emptyset] = m_{10} = L_i - L_e \quad (4.27)$$

which completes the proof of case 2.

The last case is dealt with in the same way, with the only exception that $m_4 = 0$ and instead of C_i we consider $\partial K_1 \cup \partial K_2$.

Part
Appendix

Appendix A

Review of Hexagonal Sampling *

Abstract

In this appendix we give a brief account of the state of the art in hexagonal sampling and the properties of hexagonal sampling grids, when compared to other possible regular grids in the plane.

Current imaging systems under development at the French Space Agency (CNES), which considered the use of such grids provided the initial motivation for this study. It is widely known that hexagonal sampling grids are optimal in many ways, but their use has been delayed, possibly due to the difficulty in extending efficient algorithms that have been developed for orthogonal grids. For this reason we extend our study to the algorithms that have been generalized to this case, and we give an account of the real difficulties that remain. In particular, in the case of Fourier analysis, it turns out that there exist simpler solutions than those originally considered.

The conclusions of this study provided the motivation for the research described in chapter 1 which focuses on Fourier analysis, and resolution measurement and improvement on general regular grids. For this reason some of the facts stated in this appendix are revisited with more detail in the chapter 1.

1 Introduction

Except for some experimental work with optical computers which treat analogous images by optical means and holograms, most proven image analysis technologies today resort to digital means, both for storing images and operating on them.

Thus a method has to be chosen to transform the original analogous image, into a digitally storable image. This procedure, known as digitalization, comprises a (2D) discretization of the image plane, as well as a (1D) discretization of the image intensity values. The latter problem being more simple, we shall concentrate mainly on the

*A preliminary version of this appendix was presented as my DEA thesis [6].

first problem, namely the discretization of the image plane. Among the many different ways to perform this operation, taking samples of the image from a regularly spaced square grid is the most commonly used one. A hexagonal grid, however, has many more attractive properties than the square grid, and often allows to solve problems more effectively and efficiently, as we shall point out throughout this review.

This appendix is organized as follows. Section 2 discusses the problem of digitalization and characterizes the different options we have, according to a certain set of assumptions. As a result, the square and hexagonal grids are singled out as the only reasonable possibilities. Then section 3 discusses the relative merits of both methods from a theoretical viewpoint, whereas sections 4 and 5 show how these properties become important when implementing respectively morphological and linear operators. In particular, section 6 presents some methods to construct wavelets on a hexagonal grid, and section 7 discusses how other miscellaneous operations can be implemented in both cases.

It should be emphasized that most of this appendix is organized as a review of the state of the art in hexagonal image processing, and contains little original material. It was the starting point of this thesis in 1999, and the bibliography has not been thoroughly updated since then, except for a few isolated and remarkable cases. In particular, a similar review by Staunton [168], appeared during the later stages of this work. It puts a stronger emphasis on the hardware needed for hexagonal image processing, whereas this review concentrates mainly on the mathematical and algorithmic aspects. All in all there might be a certain overlap between both reviews. Finally in appendix A and section 5.4 we discuss the accuracy and efficiency trade-offs that we face when converting between square and hexagonal grids, and we described a new algorithm we developed to solve them.

2 Digitalization

Digitalizing an image means partitioning the image plane into regions or tiles, and then assigning to that region a value that represents the light intensity of the analog image within that region, to conform a “pixel” or picture element.

This section explores the many ways in which such a partition can be made. To start, let us observe that the feasibility of an image acquisition system, already imposes a restriction on how such a partition can be made. In the particular case of satellite images, for instance, it consists of a sliding bar of equally spaced sensors, as described

in [147]. The relative orientations of the bar of sensors and the motion, as well as the sampling rate can still be freely chosen.

This imposes several constraints on the possible tilings, namely:

1. The fact that the sensors are all equal imposes a constraint on the shape of the tiles, which should resemble the region over which the sensor captures information. Hence all tiles should have the same shape.
2. Secondly, the fact that the sensors are equally spaced within the bar (and assuming a constant sampling rate) means that the centers of the tiles should lie on the lattice $n_1\mathbf{v}_1 + n_2\mathbf{v}_2$, where $(n_1, n_2) \in \mathbb{Z}^2$ are integers, $\mathbf{v}_1 \in \mathbb{R}^2$ determines the orientation of the bar of sensors and the separation between them, and $\mathbf{v}_2 \in \mathbb{R}^2$ determines the direction of motion, as well as the constant sampling rate.
3. Finally, to avoid redundancy, we should impose that the plane is a disjoint union of the tiles, which further restricts the possible shapes of the tiles.

These assumptions, along with further considerations on the spectral contents of the original and digitized images, leaves only a few choices for reasonable tilings [116, 149]. In the following section, however, we shall slightly relax the previous assumptions. This will give us more insight on the properties of the different tilings, and will enable us to arrive to the same conclusion, from very different considerations.

2.1 Tilings of the plane

To study the different ways in which the plane can be partitioned, we shall slightly modify conditions 1-3, in the following manner:

1. As in the previous section we shall consider all tiles to have the same shape $T \subset \mathbb{R}^2$. In addition we shall assume that T is a compact connected set, which is quite reasonable in view of the previous discussion.
2. Instead of requiring the centers of the tiles to lie on a lattice (or put another way, that they are obtained by translations in two directions), we shall consider a subgroup G of the group $\text{Is}^+(\mathbb{R}^2)$ of rigid transforms (or direct isometries) of the plane. This subgroup may also contain rotations, in addition to translations.

3. Finally the disjoint union condition means in this setting that:

$$\bigcup_{g \in G} g(T) = \mathbb{R}^2 \quad (\text{A.1})$$

$$\text{if } g(\overset{\circ}{P}) \cap h(\overset{\circ}{P}) \neq \emptyset \text{ then } g(P) = h(P) \quad (\text{A.2})$$

Berger [29, section 1.7] characterizes all possible tiling groups G satisfying equations (A.1) and (A.2). In fact, he shows that G is always discrete (generated by a finite number of elements). Furthermore, he shows, the subgroup $\Gamma = G \cap Tr(\mathbb{R}^2)$ of translations contained in G is actually a *lattice*, *i.e.* there is a basis $\{\mathbf{v}_1, \mathbf{v}_2\}$ of \mathbb{R}^2 such that Γ is exactly the set of translations with vectors in $\mathbb{Z}\mathbf{v}_1 + \mathbb{Z}\mathbf{v}_2$.

Since the translation part of G has always the same structure, the different tilings will be distinguished by the rotations they contain. In fact there are only 5 cases as shown in figure A.1. The first figure (a) corresponds to the case when there is no rotation, figure (b) to the case when there is a rotation of order 2 (and none of higher order, and figures (c), (d) and (f) to the cases when there is a rotation of order higher than 2. In such a case, G must contain three different rotations, of angles $2\pi/\alpha_1$, $2\pi/\alpha_2$ and $2\pi/\alpha_3$, where $\alpha_i \in \mathbb{N}$, for $i = 1, 2, 3$, $\alpha_i \geq 3$ for at least one $i \in \{1, 2, 3\}$, and $\sum \frac{1}{\alpha_i} = 1$. This means that there are only three cases as shown in the table below:

	$1/\alpha_1$	$1/\alpha_2$	$1/\alpha_3$
Case (c)	1/3	1/3	1/3
Case (d)	1/2	1/4	1/4
Case (e)	1/2	1/3	1/6

If in addition to direct isometries, we consider subgroups of $Is(\mathbb{R}^2)$, *i.e.* all isometries, including reflections, we find 12 other tilings. The corresponding drawings can be found in [29, pp. 39–40]. Another representation of the 5+12 tiling groups, which explicitly shows the centers of rotations and the axes of reflections and glide-reflections is shown in [90, pp 40–42].

We should note two more things. First the shapes of the tiles in figure A.1 can be modified in many different ways, which are compatible with the group G . Figure A.2 gives an example for case (a).

Secondly, figure A.1 shows marked tiles, in such a way that $g(\overset{\circ}{P}) \cap h(\overset{\circ}{P}) \neq \emptyset$ actually means that $g = h$ rather than $g(T) = h(T)$. We can also relax this and choose the tiles to be invariant under the action of some or all members of $G - \Gamma$. Applying this to cases (c), (d) and (e), we obtain the triangular, square and hexagonal tilings shown in

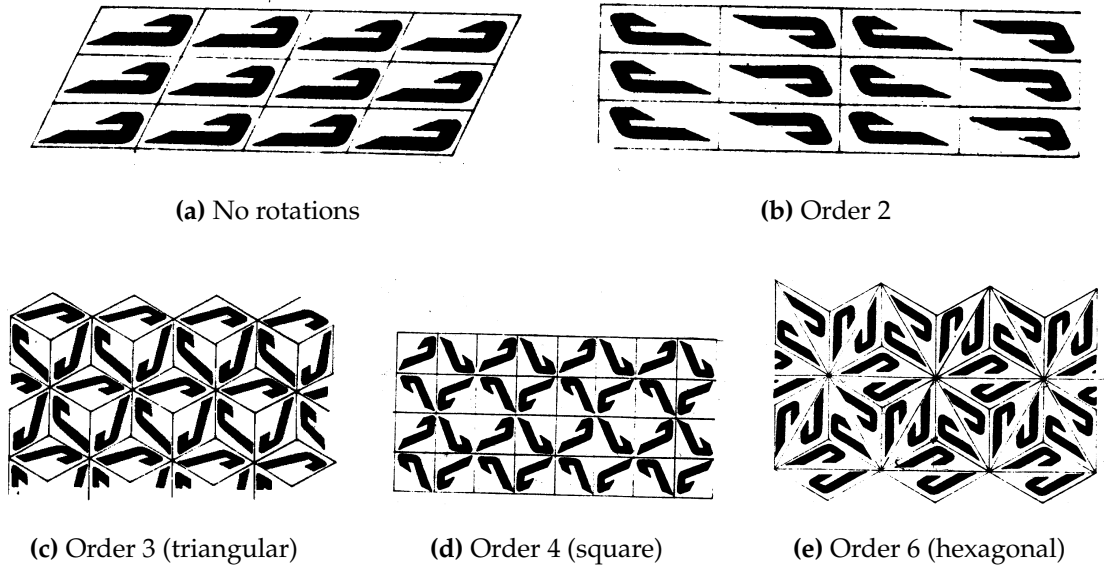


Figure A.1: The five possible tiling groups by direct isometries. G contains: (a) no rotations; (b) one rotation of order 2; (c) rotations up to order 4; (e) rotations up to order 6. (This figure was reproduced from [29, page 33, figures 1.7.4.1 to 1.7.4.5], which in turn was reproduced from [39]).




Figure A.2: The shape of a tile can be modified, and still admit the same symmetry group G . Thus, the two tiles above are admissible for case (a) in the previous figure. (Adapted from [29, page 34, figures 1.7.4.6 to 1.7.4.7]).

figure A.3. These are the only edge-to-edge tilings by regular polygons, and will be the subject of the next section 2.2.




Figure A.3: *The three regular tilings of the plane.* Observe that the symmetry groups for these tilings correspond respectively to figures A.1(c) (triangular), A.1(d) (square), and (c) to A.1(e) (hexagonal). The hexagonal tiling possesses the largest symmetry group. In addition only the square and hexagonal ones are tilings if we only allow integer translations in the group G . The triangular tile requires, in addition to translations, a symmetry with respect to an horizontal axis in order to cover the whole plane. For this and other reasons the triangular tiling is not very useful in image processing. (Adapted from [90, page 21, figure 1.2.1]).

2.2 Regular tilings and tilings by regular polygons

A more general and thorough classification of tilings is given in [90] . Under this general classification tilings in the sense of the previous section are called *monohedral*, meaning that they are generated by only one tile. The group G is called the *symmetry group* of the tiling, since it leaves the tiling globally invariant.

Within this class of monohedral tilings we are interested in those having the richest symmetry group. This can be achieved by letting the tiles be regular polygons and yields the three tilings shown in the previous section. The fact that these three are the only edge-to-edge tilings by regular polygons is proved in [90, section 2.1] and in [30, section 12.6]. In fact, both prove more general results: the first characterizes all polyhedral edge-to-edge tilings of the plane by regular polygons, while the latter characterizes all monohedral tilings of \mathbb{R}^n by regular polytopes.

Another concept which also leads to the same three tilings is the concept of regularity. Any tiling is said to be *regular* if its symmetry group is transitive on the flags. By *flag* in a tiling we mean all triplets (V, E, T) , consisting of a vertex V , an edge E and a tile T which are mutually incident. It can be shown that a regular tiling consists of regular polygons, and thus this characterization leads to the same three tilings in figure A.3.

Again [90] gives many examples of the tilings that are obtained if we impose transitivity of the symmetry group only on the tiles (isogonal tilings), edges (isotoxal tilings), or vertices (isohedral tilings); these pictures help understand why we need to impose transitivity on the flags to obtain the regularity we want.

Finally let us make an observation which distinguishes the triangular tiling from the square and hexagonal ones. Unlike hexagonal and square tilings, the triangular tiling is not generated by a subgroup consisting only of translations (*i.e.* Γ). In fact, in the square and hexagonal tilings, for all non-translational transformations $g \in (G - \Gamma)$, there is a translation in Γ having the same effect on the grid. But this is not the case for the only rotation of the triangular tiling. More intuitively, the triangular grid contains both upward and downward pointing triangles, and only one of them can be obtained from the original prototile by mere translations.

Hence, square and hexagonal grids can be generated from a prototile by mere translations (*i.e.* they satisfy equations (A.1) and (A.2) for $G = \Gamma$), whereas the triangular grid cannot. Put another way, the triangular grid fails to satisfy condition 2, as stated at the beginning of this section, or using the terminology of Serra [159] it doesn't have a module structure.

2.3 Tiling-grid duality

If for each tile in any of the regular tilings in figure A.3 we consider the center of mass, then the edges of the tiling clearly form the Voronoi diagram of the centers. Furthermore if we join by a line segment the centers of every two tiles which share an edge, we obtain a grid which is the dual of the tiling. It is interesting to note that if we consider the edges of this grid as delimiting a tile, then the dual grid of a triangular tiling is hexagonal, and vice-versa, whereas the square grid is self-dual, except for a translation of half its diagonal (see figure A.4). In the particular case of the hexagonal tiling, the resulting triangular grid is the Delaunay triangulation of the centers of the hexagons.




Figure A.4: *Tiling-grid duality* By taking the centers of the tiles, and joining the centers of adjacent tiles by edges, we obtain a new (dual) tiling. Triangular and hexagonal tilings are dual, whereas the square tiling is self-dual. (This figure is reproduced from [43, page 32, figure 1.1.3]).

2.4 Conclusion

We have analyzed different ways to tile the image plane, under a set of reasonable assumptions. In all cases the analysis leads to three possibilities, namely the triangular, square and hexagonal tilings. Within these three, the hexagonal is singled out as the one having the richest symmetry group (6 symmetries and rotations up to order 6, against 4 and 3 for the square and triangular tilings). Furthermore, the triangular tiling (unlike the other two) fails to have a module structure, and should be discarded. Maybe the only reason to keep it is that it represents the grid which is dual to the hexagonal lattice. This property, however, is unlikely to be required in image processing applications.

3 Properties of Square and Hexagonal Grids

In this section we explore the properties of the square and hexagonal grids, which according to the previous section are the only possible ones under reasonable assumptions. Occasionally we will show some results also for the triangular grid, which may be useful as well.

We know from the previous sections that (in the first two cases) the centers of the tiles will lie on a grid $\Gamma = \mathbb{Z}\mathbf{v}_1 + \mathbb{Z}\mathbf{v}_2$, where $\{\mathbf{v}_1, \mathbf{v}_2\}$ is a basis of \mathbb{R}^2 . For the square grid Γ_4 we can choose for instance $\mathbf{v}_1 = T \begin{pmatrix} 1 \\ 0 \end{pmatrix}$, $\mathbf{v}_2 = T \begin{pmatrix} 0 \\ 1 \end{pmatrix}$, whereas for the hexagonal grid

Γ_6 , we can choose $\mathbf{v}_1 = T \begin{pmatrix} 1 \\ 0 \end{pmatrix}$, $\mathbf{v}_2 = T \begin{pmatrix} 1/2 \\ \sqrt{3}/2 \end{pmatrix}$, where T is the spacing between adjacent points in the grid.

The triangular grid, however, has no module structure, but it can be considered as a sort of quincunx grid: $\Gamma_3 = \Gamma'_3 \cup \Gamma''_3$, where $\Gamma'_3 = \mathbb{Z}\mathbf{v}_1 + \mathbb{Z}\mathbf{v}_2$, $\Gamma''_3 = \frac{2}{3}(\mathbf{v}_1 + \mathbf{v}_2) + \mathbb{Z}\mathbf{v}_1 + \mathbb{Z}\mathbf{v}_2$, with $\{\mathbf{v}_1, \mathbf{v}_2\}$ as for the hexagonal grid. This should, however, not be confused with the quincunx representation of the square grid, since –unlike the square case–, in the triangular grid, the tiles corresponding to the Γ'_3 and Γ''_3 subgrids have different orientations.

Then given an analog image $f : \mathbb{R}^2 \rightarrow \mathbb{R}$, the digitalization consists of constructing a discrete image $\{y_x\}_{x \in \Gamma}$, where y_x depends primarily on the values of f within the tile centered at x . Our main concern will be of course that the digital image preserves as many as possible of the properties of the original analog image. The extent to which these properties are preserved depends of course on the procedure used to obtain y_x from the values of f within the corresponding tile, but it also depends on the choice of the grid, as it shall become apparent throughout this section. An extensive treatment of these questions can be found in [159, chapters 5-6] and [43, chapter 1], of which we present a summary here.

3.1 Discrete Connectivity

One of the properties that depends most heavily on the grid choice is the possibility to define a concept of discrete connectivity, which is coherent with the Jordan curve theorem in \mathbb{R}^2 :

Theorem 14 (Jordan Curve Theorem in \mathbb{R}^2). *Every continuous, simple, closed curve $\mathcal{C} : [0, 1] \rightarrow \mathbb{R}^2$ divides the plane in two open sets $\mathbb{R}^2 - R(\mathcal{C}) = C_1 \cup C_2$. One of these sets (the interior) is simply connected and bounded and its boundary is the range of the curve: $\partial C_1 = R(\mathcal{C})$. Hence, any continuous curve $\alpha : [0, 1] \rightarrow \mathbb{R}^2$, joining an interior point $P_1 = \alpha(0) \in C_1$ with an exterior point $P_2 = \alpha(1) \in C_2$ intersects γ in at least one point.*

In order to translate this theorem to the discrete case, we need to define the concepts of continuity and connectivity in the discrete case.

Definition 15 (B-adjacency). *Two points x and y on a grid Γ are said to be B-adjacent if $y \in \{x\} \oplus B$.*

Adjacency is symmetric and reflexive, as long as the structuring element or *neighborhood* B contains the origin, and is symmetric with respect to it.

As for the choice of B , the most common choices are the following:

In the hexagonal grid Γ_6 there is only one sensible choice, namely

$$B_6 = \{0, \mathbf{v}_1, \mathbf{v}_2, -\mathbf{v}_1, -\mathbf{v}_2, \mathbf{v}_1 - \mathbf{v}_2, -\mathbf{v}_1 + \mathbf{v}_2\}.$$

It consists of all points at distance T from the origin. It can also be defined as the set of all tiles sharing an edge with the tile in the origin.

In the square grid Γ_4 there are two possibilities:

- The 4 neighbors which are at distance T from the origin: $B_4 = \{0, \mathbf{v}_1, \mathbf{v}_2, -\mathbf{v}_1, -\mathbf{v}_2\}$. This is also the set of tiles sharing an edge with the origin.
- The 8-neighbors, which add the diagonals: $B_8 = B_4 \cup \{\mathbf{v}_1 + \mathbf{v}_2, -\mathbf{v}_1 + \mathbf{v}_2, \mathbf{v}_1 - \mathbf{v}_2, -\mathbf{v}_1 - \mathbf{v}_2\}$, *i.e.* it adds the tiles sharing one vertex with the origin

In the triangular grid Γ_3 we can define three different neighborhoods. Due to its pseudo-quincunx structure, however, the definition of these depends on whether we are on the grid Γ'_3 or Γ''_3 . We give the definitions for the Γ'_3 grid, and observe that the corresponding neighborhoods for the Γ''_3 grid consist of a horizontal reflection of them.

- B_3 is the origin plus the 3 points sharing an edge with it (all points are at distance $\frac{2}{3}T$).
- B_{12} is the origin plus the 12 points sharing a vertex with it. Within these points 3 are at distance $\frac{2}{3}T$, 6 at distance T and 3 at distance $\frac{4}{3}T$.
- B_9 consists of the same points as B_{12} except the three which are at distance $\frac{4}{3}T$.

Definition 16 (B -connectivity). *A discrete B -path (the equivalent of a continuous curve) on a grid Γ is a finite sequence $P = \{x_n\}_{i=1}^N$ of points $x_n \in \Gamma$, such that each pair of successive points is B -adjacent. Two points x and y on the grid are B -connected if there exists a B -path joining them (with $x_1 = x$ and $x_N = y$). A subset $X \subseteq \Gamma$ is B -connected if all pair of points within the set are connected by paths involving points of the set.*

The B -connectedness relationship is clearly the transitive closure of the B -adjacency relationship, and hence an equivalence relationship. The corresponding equivalence classes are called *connected components*. Observe, however, that a set $X \subseteq \Gamma$ consisting of a single connected component, may still have “holes”. To avoid this we define simple connectedness as follows:

Definition 17 (Simply connected sets). A path is called simple or an B -arc (the discrete equivalent of a simple curve) if for all $n \in \mathbb{N}$, $1 < n < N$ x_n has exactly two B -adjacent points in P and it is called a B -loop (the discrete equivalent of a closed curve) if x_1 and x_N are adjacent.

Two B -paths $P = \{x_1, \dots, x_N\}$ and $P' = \{x_1, \dots, x_k, x'_1, \dots, x'_M, x_{k+1}, \dots, x_N\}$ are B -adjacent if x'_i are all B -adjacent to either x_k or x_{k+1} . The transitive closure of the path adjacency relation is an equivalence relation between paths. A set $X \subseteq \Gamma$ is called B -simply connected, if it is B -connected and for all pairs of points in X all paths joining these two points are equivalent.

Now we are ready to state the Jordan curve theorem in the discrete case. Observe, however, that in the square grid, if we use B_4 connectedness, the complement of a simple loop can contain more than two simply connected components (see figure A.5(a)). Similarly if we use B_8 connectedness, the complement of a simple loop can consist of a single simply connected component (see figure A.5(b)). Similar situations occur in the triangular grid. On the other hand, in the hexagonal grid, the Jordan curve theorem can be translated directly:

Theorem 15 (Jordan curve for hexagonal grids). Every simple B_6 -loop L divides the hexagonal grid Γ_6 in two B_6 -connected components $\Gamma_6 - L = C_1 \cup C_2$. One of these sets (the interior) plus the loop is B_6 -simply connected and bounded and its boundary is the loop: $\partial_6(C_1 \cup L) = L$. (Here the boundary of a set X is defined as $\partial_6 X = X - (X \ominus B_6)$). Hence, any path joining points in C_1 and C_2 hits the loop L in at least one point.

In the square grid we need a turnaround to make this theorem valid, we have to use 4-connectivity for the loop and 8-connectivity for the sets (or conversely):

Theorem 16 (Jordan curve for square grids). Every simple B_4 -loop (B_8 -loop) L divides the square grid Γ_4 in two B_8 -connected components (B_4 -connected components) $\Gamma_4 - L = C_1 \cup C_2$. One of these sets (the interior) plus the loop is B_8 -simply connected (B_4 -simply connected) and bounded and its boundary is the loop: $\partial_4(C_1 \cup L) = L$ ($\partial_8(C_1 \cup L) = L$). Here the boundary of a set X is defined as $\partial_4 X = X - (X \ominus B_4)$ or $\partial_8 X = X - (X \ominus B_8)$. Hence, any B_8 -path (B_4 -path) joining points in C_1 and C_2 hits the loop L in at least one point.

A similar result is valid for the triangular grid if we use B_3 and B_{12} connectivity. It is not known to the authors whether there is an equivalent Jordan curve theorem involving B_9 connectivity in the triangular grid.

It should be noted that the use of two complementary connectivity concepts in the square or triangular grids, makes algorithms more complex, and makes it impossible to develop algorithms that operate in the same manner on an image and its complement.

Another way to turn around this problem in the square grid, is to define a B_6 neighborhood, but this is even less isotropic than the B_4 and B_8 neighborhoods normally used.

3.2 Optimal Sampling

The performance of the different grids for representing a given signal can be measured in different ways, depending on the assumptions that are made on the original image, and the properties that should be preserved in the digitized image. In most cases the hexagonal grid is optimal in terms of the number of samples that are needed to accurately represent a given signal.

Maximal Density Packings and Minimal Density Coverings


In section 2 we studied all the ways in which we can cover the plane with a disjoint union of displaced tiles, where the displacements should be within a subgroup of the rigid transforms. We concluded that only three of these tiling groups were reasonable, which in turn imposes certain restrictions on the shape of the tile.

On the other hand the acquisition system may also impose a restriction on the shape of the tile (given e.g. by the support of the MTF or the PSF), which is often incompatible with the restriction imposed by the tiling group. This happens for instance if the tile is required to be a circle.


Hence we shall weaken the disjoint union condition, which leads to the study of packings and coverings, which are lower and upper approximations to the disjoint union condition. More precisely:

Definition 18 (Packing and Covering). *Given a couple (G, T) , where the tile $T \subset \mathbb{R}^2$ is a compact subset of the plane, and $G \subseteq \text{Is}(\mathbb{R}^2)$ is a subset of rigid transforms (not necessarily a subgroup), we say that (G, T) is a covering of the plane \mathbb{R}^2 if it satisfies equation (A.1), and we say that it is a packing of the plane if it satisfies equation (A.2).*


In what follows we shall be interested in coverings and packings with circular tiles, where G consists just of translations. In fact rotations and reflections are irrelevant for a circular tile. In addition, the extent to which a covering or a packing resemble a tiling will be given by its density:



(a) 4-connectivity



(b) 8-connectivity



(c) 6-connectivity

Figure A.5: *Paradoxal situations.* (a) If we use 4-connectedness in the regular grid, then a simple closed curve can delimit more than 2 connected components. (b) if we use 8-connectedness, the complement of a simple closed curve can be a single connected component (the "interior" and the "exterior" of the curve can connect through the curve). (c) using the natural 6-connectedness of the hexagonal grid, a simple closed curve, always divides the planes in 2 regions.

Definition 19 (Density of packing or covering). The density of a covering or packing (G, T) is given by $\rho(G, T) = \lim_{s \rightarrow \infty} (n\mu(T)/s^2)$, where $\mu(T)$ is the Lebesgue measure (the area) of the tile, and n is the number of elements $g(T)$, $g \in G$ of the covering (packing) which are included in a square of area s^2 which is centered at the origin.

It is clear from the definition that coverings have $\rho \geq 1$, packings have $\rho \leq 1$ and tilings (which are both packings and coverings) have density 1. Looking at figure A.6, we see that the computation of the packing density for circular tiles can be reduced to computing the ratio between the area of a circle and its circumscribed regular polygon (triangle, square or hexagon, depending on the grid). Similarly the covering density is reduced to area ratio between a circle and its inscribed regular polygon. Table A.1 shows the results of these calculations (cf. [136, 168]).




Figure A.6: Densest packings of circular tiles for the square and hexagonal grids.

Grid	Packing Density	Covering Density
Triangular	$\frac{\pi}{3\sqrt{3}} \approx 60.46\%$	$\frac{4\pi}{3\sqrt{3}} \approx 241.84\%$
Square	$\frac{\pi}{4} \approx 78.54\%$	$\frac{\pi}{2} \approx 157.08\%$
Hexagonal	$\frac{\pi}{2\sqrt{3}} \approx 90.69\%$	$\frac{2\pi}{3\sqrt{3}} \approx 120.92\%$

Table A.1: Maximal (minimal) densities of packings (coverings) by circles for the three regular grids.

Wütrich [174] and the references therein [27, 143] show that among all coverings of the plane (G, T) by circles T , the minimum density $\rho(G, T)$ is attained when G is the group which generates the hexagonal lattice. Similarly among all packings of the plane (G, T) by circles T , the maximum density is attained when G is the group which generates the hexagonal lattice. The proof seems to date back to 1940 by Fejes Tóth [73], see also [50, p. 8-9].

This is a strong result, since it shows that if the sensors have circular symmetry, the hexagonal grid is optimal, not only among the regular grids, but also among all possible dispositions of the sensors on the plane!

As a consequence, the minimum number of samples that are needed to discretize an image with a certain regularity will be attained when the sampling grid is hexagonal. The next sections show this in the cases where the regularity of the image is measured by its spectral contents (section 3.2) or by its morphology (section 3.2).

Shannon sampling theorem

Most often the regularity of a signal is measured in terms of the support of its Fourier transform, or by the speed with which it decreases at infinity. Assume that an analog image f is known to be *band-limited*, meaning that its Fourier transform \hat{f} is zero (or negligible with respect to noise) outside of a compact set $S \subset \mathbb{R}^2$. As observed by Rougé [149], this is often the case when we know the MTF of the acquisition system. In such a case, Shannon's theorem states that the signal can be completely reconstructed by its samples over a regular grid Γ as long as this grid admits a reciprocal cell $D^* \supseteq S$ containing the support of \hat{f} (see section 5.1 or [135, 149] for a definition of reciprocal cell and the notation used in this section). The reconstruction is achieved by convolution with the inverse Fourier transform of the indicator function of D^* :

$$\tilde{f} = \Delta_{\Gamma} \cdot f \quad (\text{A.3})$$

$$f = \frac{1}{|D^*|} \overline{\mathcal{F}}(\mathbb{I}_{D^*}) * \tilde{f} \quad (\text{A.4})$$

or equivalently

$$f(x) = \frac{1}{|D^*|} \sum_{\gamma \in \Gamma} f(\gamma) \overline{\mathcal{F}}(\mathbb{I}_{D^*})(x - \gamma) \quad (\text{A.5})$$

Shannon's condition is in fact a packing condition, as defined in the previous section. Since the sampling resolution is directly proportional to the size of the reciprocal cell D^* , the optimal sampling grid will be the one with the smallest reciprocal cell which contains S . But this is equivalent to say that (Γ^*, S) is a packing with the largest possible density $\rho = |S|/|D^*|$.

When S is a circle (as is often the case, at least approximately), then we know from the previous section that the optimal sampling is achieved when the dual grid Γ^* is hexagonal, which means that Γ is also hexagonal. In addition the results obtained in the previous section for packing densities mean that when a signal is circularly band-limited, critical hexagonal sampling requires $1 - \frac{78.54}{90.69} = 13.4\%$ less samples than critical square sampling.

When S is not exactly circular, it may not be obvious how to choose the shape of the reciprocal cell D^* in such a way that it is admissible for a given grid geometry, and meets Shannon's criterion. Klaas [111] gives a characterization, which describes all admissible reciprocal cells for the hexagonal grid, which are in addition symmetrical with respect to the horizontal and vertical axes. Figure A.7 reproduces some examples shown in his paper. Observe in particular how the reciprocal cell in subfigure e closely approximates the shape of the essential support of the MTF for SPOT 5 as shown in [116].




Figure A.7: Some of the reciprocal cells which are admissible for the hexagonal grid, and are in addition symmetric with respect to the horizontal and vertical axes. As it can be seen, the reciprocal cell is determined by a continuous curve which joins the horizontal axis with the midpoint between the origin and its closest north-east point in the grid. This curve must in addition stay within a rectangle contained in the first quadrant. (Reproduced from [111, page 59, Bilder 3a-e]).

Morphological sampling theorem

The Shannon theorem is suitable for discretizing images that will be later treated by linear filters, since it ensures that the filters we apply to the discretized image are coherent with the corresponding filters on the original image and the sampling process.

Similarly, when working with morphological filters, it is most convenient to sample the image in such a way that the discrete image preserves homotopy of the original image.¹ In addition we will ask that the result of applying a discrete morphological filter to the discrete image is homotopical to applying the corresponding continuous filter to the original image. To make this possible with a finite number of samples, we need to impose a certain regularity on the original image. The results are presented, in most of the literature, for binary images.

Definition 20 (Regular model). *A binary image $X \subseteq \mathbb{R}^2$ is said to be B -regular (where $B \subseteq \mathbb{R}^2$ is a structuring element), if it is open and closed with B , i.e. if $X = X \bullet B = X \circ B$.*

Theorem 17 (Morphological critical sampling). *Let $X \subseteq \mathbb{R}^2$ be an image, and $\sigma(X, a\Gamma)$ its sampling on a grid $a\Gamma$, which is assumed to be either the hexagonal or the square grid with B_6 connectivity and spacing a between neighboring tiles. Let D_r be the disk of radius $r = f(a\Gamma)$. Then the following holds:*

- (i) *if X is D_r -regular then X and $\sigma(X, \Gamma)$ are homotopic*
- (ii) *$r = f(a\Gamma)$ is the smallest radius such that (i) holds.*

This means that for a given grid resolution $a\Gamma$, we can safely sample D_r -regular images if and only if $f(a\Gamma) \leq r$. Under these conditions the original image can be reconstructed by dilation of the sampled image. For less regular images, the digital image may not be homotopical to the original image and reconstruction may be impossible.

With respect to the function f which actually measures the critical sampling, different authors have proposed (and proved) different results. Serra [159] proposed a conservative value, which coincides with the results obtained by Haralick [93] under slightly different assumptions. Later Florencio [75] claimed (and proved) that this value was actually between $\sqrt{3}$ and 2 times smaller. Table A.2 shows the results. The critical radius of regularity is given as a function of the spacing a of the sampling. We also give the value of r^2 as a function of the area A of the unit tile (observe that $A = a^2$ for the square grid, whereas $A = \sqrt{3}a^2/2$ for the hexagonal grid). Then we can compute the critical sampling resolution $1/A$ directly as a function of the degree of regularity r .

¹Two sets $A \subseteq \mathbb{R}^2$ and $B \subseteq \mathbb{R}^2$ are said homotopic if there is a one to one correspondence between $\mathcal{A} = \{\text{connected components of } A \text{ and } A^c\}$ and $\mathcal{B} = \{\text{connected components of } B \text{ and } B^c\}$, which also preserves connectivity of arbitrary unions of the elements of \mathcal{A} and \mathcal{B} . A more detailed definition of homotopy and some examples can be found in [75, 159].

Grid	(cf. Serra)		(cf. Florencio)	
	$r = f(a\Gamma)$	r^2	$r = f(a\Gamma)$	r^2
$a\Gamma_4$	$\sqrt{2}a$	$2A$	$\frac{1}{\sqrt{2}}a$	$\frac{1}{2}A$
$a\Gamma_6$	a	$\frac{2}{\sqrt{3}}A$	$\frac{1}{\sqrt{3}}a$	$\frac{2}{3\sqrt{3}}A$

Table A.2: Smallest radius of regularity r for critical morphological sampling on a grid with spacing a and tile of area A .

Whichever result is true we should note that for a given D_r -regular signal, the hexagonal lattice requires between $1/\sqrt{3}$ and $4/3\sqrt{3}$ times the samples needed to represent the same signal with a square grid preserving homotopy, *i.e.* the critical resolution for the hexagonal grid is between 42.26% and 23.03% smaller than that of the square grid.

In addition, the square grid is assumed to possess 6-way connectivity to prove this result, and we have seen that it is impossible to obtain 6-way connectivity in the square grid in an isotropic manner. Hence from the morphological point of view there are very strong arguments in favor of the hexagonal grid.

Mean Quantization Error

If we assume the original image to consist of random points which are uniformly distributed in the plane, the mean absolute error due to nearest-neighbor spatial quantization can be computed for the different grids. In [105, 106] this error was found for the hexagonal grid to be between 10% lower and 5% higher than the square grid. For the triangular grid, however, the quantization error is between 1% lower and 13% higher when compared to the square grid. It should be noted that in their calculations, a pseudo-hexagonal grid was used, having the same horizontal and vertical spacing as the square grid (*i.e.* the one generated by the vectors $\mathbf{v}_1 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$, $\mathbf{v}_2 = \begin{pmatrix} 1/2 \\ 1 \end{pmatrix}$).

Experimental studies

In [117] triangular, square and hexagonal grids are compared for a given resolution. The comparison is made in terms of the errors in peak height and peak density that occur when a surface containing randomly located peaks is digitized at a certain resolution. It turns out that the hexagonal grid gives the smallest number of errors.

Summary

We have seen under different regularity models, that for a given regularity of the image, critical sampling (*e.g.* allowing exact reconstruction) requires a smaller number of samples in the hexagonal than in the square and triangular grids. The difference varies between 13% and 42% (depending on the hypothesis) less samples than the square grid, and is much larger when compared with the triangular grid.

Conversely, for a given regularity of the original signal and sub-critical grid resolution the hexagonally sampled signal is usually closer to the original than the squarely sampled one, according to both theoretical and experimental results.

Finally, the hexagonal grid is singled out as optimal among all possible coverings and packings by circular tiles.

3.3 Recursivity

Tile-Recursivity Chassery [43] points out as a disadvantage, that the hexagonal grid is not *recursive*, in the sense that there is no way to combine a set of hexagonal tiles so that their union is also a hexagon. In the square and triangular grids, we can combine 2^2 , 3^2 , 4^2 , 5^2 , etc. cells to form a larger cell with the same geometry. Such constructions are useful for developing hierarchical algorithms like those based on quad-trees.

Nevertheless, in the hexagonal grid, there is one way to combine 7, 7^2 , 7^3 , etc. tiles so that they form an approximately hexagonal shape, which has as well 6 coarse scale tiles as neighbors. The problem with this construction (in addition to the fact that it is an approximation), is that the coarse tile (or rather the basis of the coarse grid) is rotated 19.1° with respect to the original grid. (In [172] such a construction is used to build an orthogonal wavelet basis for the hexagonal grid).

A similar problem occurs with the triangular grid, where the coarse tile is reflected horizontally with respect to its central fine-scale tile. Figure A.8

Lattice-Recursivity For most applications, however, we can relax the recursivity condition of the previous paragraph, by only requiring the coarse lattice to be a subset of the original lattice. This is always possible, without changing the geometry of the grid if the grid has a module structure (as in the square and hexagonal case). We only need to multiply the grid basis by an integer factor. Such a construction, with a factor of 2 was used by several authors [49, 97, 130, 157, 162] to construct wavelet bases of different types.

In addition we can obtain a lattice hierarchy, by selecting as the basis of the coarse




Figure A.8: Tile-Recursivity. In the triangular and square grids we can combine n^2 tiles to form a larger tile of the same shape. In the hexagonal grid, this is only possible in an approximate manner, rotating the grid basis, and increasing the side of the hexagons by $\sqrt{7}$. (Subfigure (c) was reproduced from [43, page 128, figure 4.27]).

grid any pair of linearly independent vectors $\mathbf{v}'_1, \mathbf{v}'_2 \in \Gamma$, such that $\|\mathbf{v}'_1\| = \|\mathbf{v}'_2\|$ and $\langle \mathbf{v}'_1/\|\mathbf{v}'_1\|, \mathbf{v}'_2/\|\mathbf{v}'_2\| \rangle = \langle \mathbf{v}_1/\|\mathbf{v}_1\|, \mathbf{v}_2/\|\mathbf{v}_2\| \rangle$. This includes the hexagonal case depicted in figure A.8(c).

3.4 Miscellaneous properties

Human visual system As an application of the previous ideas let us observe that many authors [168, 172] have pointed out the following fact about the human visual system. The roughly circular sensory cells in the human retina tend to pack themselves efficiently, thus leading to a hexagonal configuration. In addition, the distribution of sensory cells is roughly pyramidal, with the finest resolution near the fovea, and coarser resolution at more distant points.

Principal Directions. Serra [159] shows the following result about grids. Two non-parallel straight lines in a grid may have an empty intersection. Consider for instance the lines $\{k(\mathbf{e}_1 + \mathbf{e}_2) : k \in \mathbb{Z}\}$ and $\{\mathbf{e}_1 + k(\mathbf{e}_2 - \mathbf{e}_1) : k \in \mathbb{Z}\}$ in the square grid, which do not intersect. In that case we say that they *cross*. Otherwise, if there is a single point where they intersect we say that they *meet*.

Let Γ be a grid and α, β two primitive vectors in the grid. Then any lines $D_2 \subset \Gamma$ parallel to α meets any other line $D_2 \subset \Gamma$ parallel to β if and only if $\{\alpha, \beta\}$ define a basis of Γ .

Now let α, β, γ be three primitive vectors of Γ and u, v, w any three points of Γ where we draw lines parallel to α, β, γ respectively. Then any such line meets the other two if and only if:

1. $\{\alpha, \beta\}$ is a basis of Γ
2. $\{\alpha + \varepsilon\beta + \varepsilon'\gamma = 0$ for some $\varepsilon, \varepsilon' \in \{-1, +1\}$

In such a case we say that $\{\alpha, \beta, \gamma\}$ is a set of *principal directions*. We can verify that there can not be more than three principal directions.

This justifies the choice of the third vector in the Her coordinate system for the Hexagonal grid: it is chosen in such a way that $\{\tilde{e}_1, \tilde{e}_2, \tilde{e}_3\}$ is a set of principal directions. Otherwise we would not be able to convolve along the third direction as well. From these results it becomes also clear why we can consider 3-separable filters but not 4-separable filters in 2-D grids. In addition it becomes clear that only in the regular hexagonal grid we can have an “isotropic” set of three principal directions (with the same norm and forming the same angle to each other). In the square grid this is not possible.

Perceptual Quality Square sampling avoids jagged vertical or horizontal edges, which are perceptually very bad [140, 167]. However, this is gained only at the expense of a more discontinuous representation of edges or lines in other directions [174].

4 Morphological and level-set operations

As noted in the previous sections, the hexagonal grid has from the geometrical point of view several advantages over the square grid:

- *Connectivity* is easier to define in a way which is consistent with the continuous case. In contrast to the square case only one structuring element can be used to determine connectedness, whereas in the square case foreground and background have to use different structuring elements to be consistent with Jordan curve’s theorem.

This results in a fundamental impossibility to construct algorithms which operate in the same manner for an image and its complement. This is not the case in the hexagonal case and algorithms are largely simplified by using only one connectivity concept. We can take advantage of this fact to implement improved

version of the Fast Level Lines Transform [134], and of several algorithms based on it, *e.g.* restoration with the grain filter [123].

- The *symmetry group* of the grid is larger, meaning that it is more “*isotropic*”. Put another way, a circular structuring element of unit radius is more closely approximated in the hexagonal grid (by a hexagon), than in the square grid (by a square or a cross).

This means that the discretization of algorithms which are based on the iteration of circular structuring elements, will have better approximations to the continuous case, if we use a hexagonal grid, than if we use a square grid.

Direct and more elaborate consequences of these simple facts were extensively used in the literature. In the following sections we give some examples.

4.1 Thinning

In [167] iterative thinning algorithms on square and hexagonal grids were compared. The hexagonal version was found to better approximate the medial-axis definition of skeleton, and was found to be less sensitive to noise in the boundary of shapes. In addition the hexagonal algorithm is simpler and uses a smaller number of hit-miss templates, which also leads to improved computational efficiency.

4.2 Distance Transforms

A discrete distance measure can be defined [119] as the number of steps in the shortest path between two points. Then the distance measure depends on the connectivity of the lattice. For the square grid, using 4-neighbor connectivity leads to the L^1 -norm, whereas using 8-neighbor connectivity leads to the L^∞ -norm.

For the hexagonal grid, the 6-neighbor distance does not correspond exactly to any L^p -norm, but it more closely approximates the Euclidean distance than either the L^1 -norm or the L^∞ -norm, and it can be computed by a simple formula in the hexagonal grid.

In [38] the construction of distance transforms in the hexagonal grid is discussed. A distance transform computes the Euclidean distance from each point to the foreground of a binary image. Distance transforms have been extensively used for a variety of binary image operations, and more recently for the implementation of level-set methods for solving PDEs. Efficient (though approximate or pseudo-euclidean) implementation of a distance transform requires the propagation of global distances from local operations, as well as integer arithmetic if possible.

The simplest distance transform is given by the 6-neighbor distance, with a maximum error of 13.4%, which is still smaller than the error obtained for the 4-neighbor or 8-neighbor distance. With a larger computational cost, [38] constructs integer distance transforms with a maximum error of 3.3% and 1.0%, and real distance transforms with a maximum error of 1.6% and 0.7%.

5 Linear Filtering

Unlike the case of morphological operators, linear operators in the hexagonal grid are somewhat more difficult to implement. Fundamental problems that arise are related to the fact that isotropic filters are no longer separable in the classical sense, and to the impossibility to keep a square periodization.

Despite the higher complexity, however, better linear operators can be obtained in the hexagonal grid, at a slightly larger computational cost. Here we mean better in the sense of isotropy and angular resolution.

5.1 Review of Fourier Analysis

In this section we give a short review of the main results of Fourier theory, and we introduce the notation that we shall use to generalize the discrete Fourier transform for images that are defined in a hexagonal grid rather than the classical square grid. The presentation is based primarily on [135] and adapted to present the results of [68, 128, 157].

Recall that (in the continuous case, and under certain conditions precised in chapter 1) given a distribution f , its Fourier transform $\hat{f} = \mathcal{F}(f)$, and its inverse Fourier transform $\overline{\mathcal{F}}(\hat{f})$ can be defined by the following expressions:

$$\hat{f}(\xi) = \mathcal{F}(f)(\xi) = \int_{\mathbb{R}^2} f(x) e^{-ix \cdot \xi} dx \quad (\text{A.6})$$

$$(2\pi)^2 f(x) = \overline{\mathcal{F}}(\hat{f})(x) = \int_{\mathbb{R}^2} \hat{f}(\xi) e^{ix \cdot \xi} d\xi \quad (\text{A.7})$$

Sampling. To translate these expressions to the discrete case, we will sample the signal by means of a Dirac comb, *i.e.* a set of Dirac delta functions defined on a grid Γ :

$$\Delta_\Gamma = \sum_{\gamma \in \Gamma} \delta_\gamma \quad (\text{A.8})$$

We shall say that a function u is Γ -discrete, if it was obtained by sampling a certain function or distribution \tilde{u} , *i.e.* if:

$$u(x) = \Delta_{\Gamma} \tilde{u}(x) = \sum_{\gamma \in \Gamma} \tilde{u}(x) \delta(x - \gamma) \quad (\text{A.9})$$

To simplify the notation we shall consider the matrix $\mathbf{V} = [\mathbf{v}_1 \mathbf{v}_2]$, having as columns the two vectors forming a basis of Γ . Then any point $\gamma \in \Gamma$ in the grid can be written:

$$\gamma = \mathbf{V} \mathbf{n} = n_1 \mathbf{v}_1 + n_2 \mathbf{v}_2 \quad (\text{A.10})$$

and the Γ -discrete function u is determined by the values:

$$u[\mathbf{n}] = u[n_1, n_2] = \tilde{u}(\mathbf{V} \mathbf{n}) = \tilde{u}(\gamma) \quad (\text{A.11})$$

for $\gamma \in \Gamma$.

Periodization. In order to deal with images of finite extent, one usually assumes that the image is periodic. More precisely, we say that a distribution is Γ -periodic if it is invariant by translations within the grid, *i.e.* if

$$u(x) = u(x + \gamma), \quad \text{for all } \gamma \in \Gamma \quad (\text{A.12})$$

Equivalently u is Γ -periodic if it is the convolution of a certain distribution f with a Dirac comb, *i.e.*

$$u = \Delta_{\Gamma} * f \quad (\text{A.13})$$

For a Γ -periodic function or distribution u the signal can be reconstructed from its discrete Fourier coefficients:

$$\hat{u}[\mathbf{k}] = c_{\mathbf{k}}(u) = \frac{1}{|D|} \int_D u(x) e^{-i(\mathbf{V}^* \mathbf{k})^T x} dx \quad (\text{A.14})$$

$$u(x) = \sum_{(\mathbf{V}^* \mathbf{k})^T \in \Gamma^*} \hat{u}[\mathbf{k}] e^{i(\mathbf{V}^* \mathbf{k})^T x} \quad (\text{A.15})$$

where the equality holds provided that:

- $D \subseteq \mathbb{R}^2$ is a *unit cell* of the grid Γ , *i.e.* if (Γ, D) is a tiling of the plane. Put another way, D must be a tile, such that its translates by $\gamma \in \Gamma$ form a disjoint union of \mathbb{R}^2 .
- Γ^* is the *dual grid* of Γ , *i.e.* if the columns of \mathbf{V} form a basis of Γ , then the columns of the matrix \mathbf{V}^* such that $\mathbf{V}^T \mathbf{V}^* = (2\pi)^2 \mathbf{I}$ are a basis of Γ^* .
- $\mathbf{k} \in \mathbb{Z}^2$ and hence the frequencies $\mathbf{V}^* \mathbf{k} = \omega^* \in \Gamma^*$ belong to the dual grid.

Similarly, a unit cell D^* of the dual grid Γ^* is called a *reciprocal cell*.

Sampling and periodization in the Fourier domain. The results above can be summarized in the following terms:

If a distribution $u = \Delta_\Gamma * \tilde{u}$ is Γ -periodic its Fourier transform is Γ^* -discrete. Put another way, periodization in the spatial domain, implies sampling in the frequency domain:

$$\hat{u} = \widehat{\Delta_\Gamma} \cdot \hat{\tilde{u}} = \Delta_{\Gamma^*} \cdot |D^*| \hat{\tilde{u}} \quad (\text{A.16})$$

Conversely, when a distribution $u = \Delta_\Gamma \cdot \tilde{u}$ is Γ -discrete its Fourier transform becomes Γ^* -periodic. Put another way, sampling in the spatial domain implies periodization in the Fourier domain:

$$\hat{u} = (2\pi)^{-2} \widehat{\Delta_\Gamma} * \hat{\tilde{u}} = (2\pi)^{-2} |D^*| \Delta_{\Gamma^*} * \hat{\tilde{u}} = \Delta_{\Gamma^*} * \frac{1}{|D|} \hat{\tilde{u}} \quad (\text{A.17})$$

5.2 Discrete Fourier transform

In practice we shall consider digital images to be Γ -discrete and Γ_N -periodic. Here the Γ grid expresses the disposition of sensors, and Γ_N must be a subgrid of Γ , *i.e.*:

$$\Gamma_N = \{\mathbf{V}\mathbf{N}\mathbf{n} : \mathbf{n} \in \mathbb{Z}^2\} \quad (\text{A.18})$$

where \mathbf{N} is a 2×2 integer matrix. When Γ is a rectangular grid the common practice is to use a diagonal matrix $\mathbf{N} = \text{diag}(N_1, N_2)$, where N_1 and N_2 are the horizontal and vertical dimensions (measured in pixels) of the image. The result is then the usual rectangular or “toroidal” periodization.

When the original grid is not rectangular, however, there are other reasonable possibilities that can be considered, as we shall see below.

According to the results in the previous section, if $u = \Delta_{\Gamma_N} \cdot (\Delta_\Gamma \tilde{u})$ is Γ -discrete and Γ_N -periodic, then it is determined by the values:

$$u[\mathbf{n}] = \tilde{u}(\mathbf{V}\mathbf{n}) = \tilde{u}(\gamma) \quad (\text{A.19})$$

for $\gamma \in \Gamma \cap D$ where D is a unit cell of the Γ_N grid. The area of such a cell is $|D| = |\det \mathbf{V} \det \mathbf{N}|$ and it always contains $|\det \mathbf{N}|$ pixels, regardless of the shape we choose for D .

In addition, for a Γ -discrete and Γ_N -periodic function u , its discrete Fourier transform is Γ^* -periodic and $(\Gamma_N)^*$ -discrete, and is determined by the values:

$$\hat{u}[\mathbf{k}] = \hat{\tilde{u}}(2\pi \mathbf{V}^{-\text{T}} \mathbf{N}^{-\text{T}} \mathbf{k}) = \hat{\tilde{u}}(\mathbf{V}^* \mathbf{N}^{-\text{T}} \mathbf{k}) = \hat{\tilde{u}}(\omega) \quad (\text{A.20})$$

for $\omega \in (\Gamma_N)^* \cap D^*$ where D^* , is a reciprocal cell, *i.e.* the unit cell of the Γ^* grid. The area of such a cell is $|D^*| = |\det(\mathbf{V}^*)| = 2\pi/|\det(\mathbf{V})|$ and it always contains $|\det(N)|$ samples of the $(\Gamma_N)^*$ grid, regardless of the shape we choose for D^* .

It can be verified that the a discrete and periodic function u and its discrete Fourier transform \hat{u} are related by the following equations:

$$\begin{aligned}\hat{u}[\mathbf{k}] &= \frac{1}{|\det(N)|} \sum_{\gamma=\mathbf{V}\mathbf{n} \in \Gamma \cap D} u[\mathbf{n}] e^{-i(\mathbf{V}^* \mathbf{N}^{-\mathbf{T}} \mathbf{k})^{\mathbf{T}} \mathbf{V} \mathbf{n}} \\ u[\mathbf{n}] &= \sum_{\omega=\mathbf{V}^* \mathbf{N}^{-\mathbf{T}} \mathbf{k} \in (\Gamma_N)^* \cap D^*} \hat{u}[\mathbf{k}] e^{i(\mathbf{V}^* \mathbf{N}^{-\mathbf{T}} \mathbf{k})^{\mathbf{T}} \mathbf{V} \mathbf{n}}\end{aligned}$$

that can be simplified into:

$$\hat{u}[\mathbf{k}] = \frac{1}{|\det(N)|} \sum_{\mathbf{v}\mathbf{n} \in \Gamma \cap D} u[\mathbf{n}] e^{-2\pi i \mathbf{k}^{\mathbf{T}} \mathbf{N}^{-1} \mathbf{n}} \quad (\text{A.21})$$

$$u[\mathbf{n}] = \sum_{2\pi(\mathbf{V}\mathbf{N})^{-\mathbf{T}} \mathbf{k} \in (\Gamma_N)^* \cap D^*} \hat{u}[\mathbf{k}] e^{2\pi i \mathbf{k}^{\mathbf{T}} \mathbf{N}^{-1} \mathbf{n}} \quad (\text{A.22})$$

Choice of periodization. In the case when Γ is a regular square grid, as said before, D is usually a rectangle of $N_1 \times N_2$ pixels, and \mathbf{N} is a diagonal matrix.

When Γ is a hexagonal grid, however, several options have been considered:

Γ -periodization. Results from considering D to be a parallelogram of $N_1 \times N_2$ pixels.

In this case the most natural choice for the periodization matrix is the diagonal

$$\mathbf{N} = \begin{pmatrix} N_1 & 0 \\ 0 & N_2 \end{pmatrix}. \quad (\text{A.23})$$

This is the approach taken in [135], and it has the advantage that the discrete Fourier transform (A.21) assumes the same form as in the square case, and no new algorithm is needed. In addition, for a regular hexagonal grid, the reciprocal grid, where the Fourier coefficients are defined, is also hexagonal, having the same geometry as Γ .

Rectangular or quasi-square periodization. Results from considering D to be a rectangle of $N_1 \times N_2$ pixels. In this case if we assume the basis of Γ to be composed of the following vectors:

$$\mathbf{v}_1 = \begin{pmatrix} d_1 \\ 0 \end{pmatrix} \quad (\text{A.24})$$

$$\mathbf{v}_2 = \begin{pmatrix} d_1/2 \\ d_2 \end{pmatrix} \quad (\text{A.25})$$

the most natural choice for the periodization matrix is

$$\mathbf{N} = \begin{pmatrix} N_1 & -N_2/2 \\ 0 & N_2 \end{pmatrix} \quad (\text{A.26})$$

In fact, for this choice of \mathbf{N} the periodization grid Γ_N has as basis $\mathbf{V}\mathbf{N} = \text{diag}(N_1d_1, N_2d_2)$, and is therefore a rectangular grid. The dual grid $(\Gamma_N)^*$ where the Fourier coefficients are defined, having basis $(2\pi)^{-2}\mathbf{V}\mathbf{N}^{-T}$ is also rectangular. In particular, both the periodization and the dual grids will become square if:

$$\frac{d_1}{d_2} = \frac{N_2}{N_1} \quad (\text{A.27})$$

In addition, the original grid Γ will be a regular hexagonal grid, only if

$$\frac{d_1}{d_2} = \frac{2}{\sqrt{3}}. \quad (\text{A.28})$$

From equations (A.28) and (A.27) it becomes clear that we cannot have exactly square periodization on a regular hexagonal grid. However, for images consisting of a large number of pixels, we can obtain a close rational approximation for the irrational number $\frac{2}{\sqrt{3}}$.

This approach was considered in [68], where it was important to use in the hexagonal grid, the same periodization as was used in a rectangular grid. We shall return to this point in section 5.4.

It should be noted that since \mathbf{N} is not diagonal, in this approach, a new algorithm has to be developed to compute the discrete Fourier transform in equation (A.21). This restricts significantly the possibility to choose an arbitrarily close approximation to the square grid, because a new algorithm is needed for each different ratio N_2/N_1 . In practice $N_2/N_1 = 8/7$ is used most often, since it approximates $2/\sqrt{3}$ to an accuracy of 1% and a fast algorithm exists for the DFT in this case.

Hexagonal periodization. Results from considering D to be an irregular hexagon with the two horizontal sides of length N_1 pixels, and the four oblique sides of length N_2 pixels. In this case the periodization matrix becomes:

$$\mathbf{N} = \begin{pmatrix} N_1 + N_2 & N_2 \\ N_2 & 2N_2 \end{pmatrix} \quad (\text{A.29})$$

The resulting periodization grid Γ_N is neither hexagonal nor square or rectangular, and the matrix is not diagonal, so a new algorithm is needed for the DFT.

Despite these disadvantages, this periodization has been widely used in the literature [138, 157] since it was first introduced by the pioneering work of Mersereau [128]. The reason was probably the fact that a 25% more efficient algorithm is known for the FFT using this periodization, than for the square FFT [128, 129]. This algorithm is usually referred to as HFFT for Hexagonal Fast Fourier Transform.

5.3 Separable Filtering

A filter $f(x_1, x_2)$ is defined to be separable, if it can be expressed as a tensor product of two functions of one-variable

$$f(x_1, x_2) = f_1(x_1)f_2(x_2) \quad (\text{A.30})$$

Such a tensor product can also be considered as a convolution of two 1-D functions:

$$f(x_1, x_2) = (f_1(x_1)\delta(x_2)) * (\delta(x_1)f_2(x_2)) \quad (\text{A.31})$$

This shows the fundamental advantage of a separable filter, namely that the 2-D convolution of a signal g with f can be expressed as two 1-D convolutions.

For discrete signals this means an important saving in terms of computational cost. In fact if the image g is a square of $n \times n$ pixels and the filter has a square support of size $m \times m$ pixels, then the 2-D convolution requires m^2n^2 flops (additions and multiplications), whereas the corresponding complexity of the separable convolution (two 1-D convolutions) is only $2mn^2$ flops. For filters having a large support the difference can be significant.

The importance of separable filters in the case of square grids, is based on the fact that one can implement rotationally symmetric (or at least approximate rotationally symmetric) filters in this manner. For example the rotationally invariant Gaussian and sinc filters are separable. Other separable filters which are not rotationally invariant are also widely used such as the box filter for instance.

Separable filters are not as attractive in the hexagonal grid, because they are not as symmetric as in the square grid. They have a certain bias towards a certain direction, which depends on the chosen basis vectors.

Nevertheless, in the hexagonal grid we can consider *3-separable* filters which can be defined in the following manner.

Let Γ be a regular hexagonal grid with basis $\{e_1, e_2\}$, and let $\{\tilde{e}_1, \tilde{e}_2, \tilde{e}_3\}$ be the corresponding Her basis (see section 7.3). The corresponding Her coordinates of a

point $\mathbf{x} = x_1\mathbf{e}_1 + x_2\mathbf{e}_2$ are $(\tilde{x}_1, \tilde{x}_2, \tilde{x}_3)$ such that $\tilde{x}_1 + \tilde{x}_2 + \tilde{x}_3 = 0$. Similarly, we can define the Her coordinates for the dual grid Γ^* .

Then a filter f is said to be 3-separable, if its Fourier transform \hat{f} can be expressed in Her coordinates as a tensor product of three 1-D functions:

$$\hat{f}(\tilde{\omega}_1, \tilde{\omega}_2, \tilde{\omega}_3) = \hat{f}_1(\tilde{\omega}_1)\hat{f}_2(\tilde{\omega}_2)\hat{f}_3(\tilde{\omega}_3) \quad (\text{A.32})$$

Unlike the classical separability, this condition does not imply that f itself is a product of three 1-D functions. However, it can still be expressed as a convolution of three 1-D functions, or alternatively as a tensor product in two directions plus a convolution in the third direction:

$$f(\tilde{x}_1, \tilde{x}_2, \tilde{x}_3) = f_1(\tilde{x}_1) * f_2(\tilde{x}_2) * f_3(\tilde{x}_3) \quad (\text{A.33})$$

$$= (f_1(\tilde{x}_1)f_2(\tilde{x}_2)) * f_3(\tilde{x}_3) \quad (\text{A.34})$$

Observe that if each f_i has length m pixels, then according to the last equality, the support of f is $3m^2$. Thus, convolutions with 3-separable filters can be computed with $3mn^2$ operations instead of the $3m^2n^2$ operations that would be needed for non-separable 2-D convolution, *i.e.* 1.5 times more than the operations needed to compute a 2-separable filter of the same size.

Some interesting examples of 3-separable filters are the circular Gaussian, and the inverse Fourier transform of the indicator function \mathbb{I}_{D^*} of D^* , where D^* is a hexagon centered in the origin (this example will be of special interest in the next section). Let us note, however, that the hexagonal box filter itself \mathbb{I}_D is not 3-separable. Nevertheless, there is a low-pass filter, which is as computationally inexpensive as the box filter, and is 3-separable and continuous, namely:

$$f(\tilde{x}_1, \tilde{x}_2, \tilde{x}_3) = \mathbb{I}_A(\tilde{x}_1) * \mathbb{I}_A(\tilde{x}_2) * \mathbb{I}_A(\tilde{x}_3) \quad (\text{A.35})$$

$$(\text{A.36})$$

where $A = \{x \in \mathbb{R} : |x| \leq a\}$, and \mathbb{I}_A is the indicator function of the set A . After some calculations we can see that this is a pyramid with vertex at the origin, and hexagonal basis (see figure A.9). As we shall see in section 6.2, this filter was used by Meyer and Jaffard to construct an orthonormal wavelet basis on the hexagonal grid.

5.4 Conversion between Square and Hexagonal

Despite the theoretical advantages offered by hexagonal grids, most available images have been digitized on a square grid. In order to take advantage of hexagonal algorithms a means for converting to a hexagonal grid is necessary.




Figure A.9: An example of a 3-separable filter in the spatial domain.

Conversely, if our sensor directly digitizes in the hexagonal grid, a conversion to square grid will most often be necessary for display purposes.

Nearest neighbor interpolation. In the case when the square to hexagonal conversion is used to apply morphological filters, Serra and Laÿ propose in [158] to use nearest neighbor interpolation as a way to preserve the level sets of the original. They also study the resolution levels that are needed to preserve the homotopy of the original, as well as other less demanding properties.

Spline interpolation. Starting from a rather empirical point of view, Her and Yuan [96] analyze the performance of different interpolation methods (Nearest Neighbor, Linear, Cubic and B-Spline) in terms of the visual appearance of the converted images.

Sinc interpolation. When the MTF of the acquisition system is known in detail, or when linear filtering is going to be applied after the conversion, it is more advisable to perform some kind of sinc interpolation which according to Shannon's theorem allows to reconstruct the original image exactly, by means of equation (A.4).

Let Γ^{sq} be the original square grid with basis $\{\mathbf{u}_1, \mathbf{u}_2\}$ and Γ^{hex} be the target hexagonal grid with basis $\{\mathbf{e}_1, \mathbf{e}_2\}$. Let also D_{sq}^* be the square reciprocal cell of Γ^{sq} and D_{hex}^* be the hexagonal reciprocal cell of Γ^{hex} . Assuming that the square image was correctly sampled (*i.e.* $D_{\text{sq}}^* \supseteq S$) the dimensions of the hexagonal grid have to be chosen in such a way that $D_{\text{hex}}^* \supseteq D_{\text{sq}}^*$ (if S is known precisely we can reduce the number of required samples by choosing $D_{\text{hex}}^* \supseteq S$). Finally, in order to simplify the computations the orientation of the grid is chosen in such a way that $\mathbf{e}_1 = \lambda \mathbf{u}_1$. Finally, the values of f on the new grid Γ^{hex} can be computed by means of equation (A.4), which in this case

consists of two 1-D convolutions with sinc functions.

It should be observed that in the previous conversion procedure it was assumed that the image f is Γ_N^{sq} -periodic, where the diagonal matrix N gives the size (in pixels) of the square image. If we now want to filter the image in the hexagonal grid, or even convert it back to a square grid, we will need a periodic extension of it. This periodic extension can be performed in a way which is consistent with the initial periodicity assumption only if Γ_N^{sq} is a subgrid of Γ^{hex} . Recalling the discussion in section 5.2, we observe that this is not possible, as long as both Γ^{sq} and Γ^{hex} are regular square and hexagonal grids respectively.

In view of this problem several authors [129, 162] proposed to use a pseudo-hexagonal grid, where the vertical axis is shrunk by approximately 1% with respect to the regular hexagonal grid. With this construction, the periodicity grid is also contained in the pseudo-hexagonal grid, as long as the image size is a multiple of (7, 8). The actual procedure described in [162] consists of resampling the image vertically by a factor of 7/4 using sinc interpolation, and then multiplying by $f(i, j) = 1 + (-1)^{i+j}$. The square grid is obtained by interpolating the zero-valued pixels (the centers of the “hexagons”), and then resampling vertically by a factor of 4/7.

In this appendix, based on the ideas proposed in [150] we performed the conversion by one-dimensional convolutions as explained in appendix A by assuming a periodization along the same directions as the sampling grid. This allows to preserve the regular square and hexagonal geometry in both grids. The main disadvantage of such an approach is that the converted image is not periodic with respect to a subgrid of the target grid, so once we performed the conversion, the image cannot be periodically extended in a convenient way, and we can no longer apply the FFT (which assumes periodicity with respect to a subgrid). There are two possible solutions to this problem:

1. Do all filtering in the spatial domain, and anticipate a periodical extension of the image, before conversion, which is large enough for all subsequent linear filtering operations.
2. Before conversion, multiply the image by a smooth (Hanning, Hamming or spline) window, which is equal to 1 inside the image domain, except near its borders where it smoothly tends to 0. Then we can zero-extend the image and change the periodization grid without significantly affecting the interpolation error.

6 Wavelets

6.1 Classical construction

The usual method to generalize wavelet bases to 2-D, is based on tensor products. In fact, if φ and ψ are respectively the scaling function and mother wavelet from a 1-D multiresolution analysis, then the corresponding 2-D multiresolution analysis will have a single scaling function $\varphi(x)\varphi(y)$ and three wavelets $\{\psi(x)\varphi(y), \varphi(x)\psi(y), \psi(x)\psi(y)\}$ which generate the space of details W_1 .

The main problem with wavelet bases obtained from such a tensor product construction is the ambiguity that results between the two orthogonal diagonal orientations. Whereas it is possible to distinguish vertical from horizontal frequencies from the wavelet coefficients, it is not possible to tell one diagonal orientation from the other.

As a way to avoid this problem Kovacevic, Vetterli, Karlson and Andrews [21, 22, 113] proposed to use a quincunx grid, where each subsampling step involves a rotation of 45 degrees of the square grid.

6.2 Orthogonal wavelet bases with hexagonal symmetry

Meyer and Jaffard [97, 130] propose a multiresolution analysis for constructing orthonormal wavelet bases of $L^2(\mathbb{R}^n)$ which can be defined on affine grids $\Gamma \subset \mathbb{R}^n$. In the particular case of hexagonal grids they show hexagonal symmetry. The construction starts from a function g such that $\{g(x - \gamma)\}_{\gamma \in \Gamma}$ is a Riesz basis of a space V_0 from the multiresolution analysis. Then they show that the scaling function φ having Fourier transform:

$$\hat{\varphi}(\omega) = \frac{\hat{g}(\omega)}{\sqrt{\sum_{\gamma^* \in \Gamma^*} |\hat{g}(\omega + \gamma^*)|^2}} \quad (\text{A.37})$$

yields an orthonormal basis $\{\varphi(x - \gamma)\}_{\gamma \in \Gamma}$ of V_0 , and consequently $\{2^{nj/2}\varphi(2^j x - \gamma)\}_{\gamma \in \Gamma}$ is a basis of \bigcup_{V_j} whose closure is $L^2(\mathbb{R}^n)$.

Now if we call m_0 the corresponding low-pass filter such that:

$$\hat{\varphi}(2\omega) = \hat{m}_0(\omega)\hat{\varphi}(\omega) \quad (\text{A.38})$$

then an orthonormal wavelet basis $\{2^{nj/2}\psi_r(2^j x - \gamma)\}_{\gamma \in \Gamma, j \in \mathbb{Z}, r=1, \dots, 2}$ of $L^2(\mathbb{R}^n)$ is obtained for

$$\hat{\psi}_r(2\omega) = m_r(\omega)\hat{\varphi}(\omega) \quad (\text{A.39})$$

under certain conditions for the high-pass filters m_r . For the particular case of the two-dimensional ($n = 2$) regular hexagonal grid, Meyer and Jaffard show that these filters

can be obtained by modulation of the low-pass filter as follows:

$$m_1(\omega) = e^{i\langle w, e_1 \rangle} m_0\left(w + \frac{1}{2}e_1^*\right) \quad (\text{A.40})$$

$$m_2(\omega) = e^{i\langle w, e_2 \rangle} m_0\left(w + \frac{1}{2}e_2^*\right) \quad (\text{A.41})$$

$$m_3(\omega) = e^{i\langle w, e_3 \rangle} m_0\left(w + \frac{1}{2}e_3^*\right) \quad (\text{A.42})$$

$$(\text{A.43})$$

where $\{e_1, e_2\}$ is the basis of the grid Γ , $e_3 = e_2 - e_1$, and $\{e_1^*, e_2^*, e_3^*\}$ is 2π times $\{e_1, e_2, e_3\}$ rotated by $\frac{\pi}{6}$. If we assume that the low-pass filter has essentially hexagonal support, then this multiresolution analysis can be illustrated in the Fourier domain by a diagram like figure A.10.

As an example of this construction, Meyer and Jaffard consider the space V_0 consisting of piecewise affine functions on the hexagonal grid, which admits spline functions as a Riesz basis. Then, to obtain an orthogonal wavelet basis with an arbitrary degree of regularity, they propose to do the same construction with $G = g * \cdots * g * g$ instead of g .

6.3 Hexagonal subband coding

Simoncelli [161, 162], uses a similar construction to produce a perfect reconstruction subband coding scheme with hexagonally symmetric filters of small support. However, instead of starting from a Riesz basis, he constructs directly a low pass filter m_0 and the corresponding high pass filters, according to equation (A.40), with the only condition that:

$$\sum_{k=0}^3 \left| m_0\left(w + \frac{1}{2}e_k^*\right) \right|^2 = 4 \quad (\text{A.44})$$

Actually the construction is done by fixing the size of the filter, imposing the hexagonal symmetry constraints, and then computing the coefficients of \hat{m}_0 by minimizing the “perceptual aliasing error”:

$$E = \max_w \frac{1}{|w|} \left| 4 - \sum_{k=0}^3 \left| m_0\left(w + \frac{1}{2}e_k^*\right) \right|^2 \right|$$

With respect to the previous construction, this one raises the following questions:

- Does the construction lead to orthogonal wavelet bases of V_0 and $L^2(\mathbb{R}^2)$? In this respect A. Cohen [48, 49] gives certain necessary and sufficient conditions that can be used to verify this.




Figure A.10: Idealized representation of the multiscale analysis resulting from the hexagonal QMFs proposed in [162]. Reproduced from [162, page 175, figure 416].

- Are the filters 3-separable? The filters are designed to have a small support, and exploit hexagonal symmetry, so that a small number of multiplication (but a large number of additions) is needed for the convolution. 3-separable convolution would be of interest to reduce both the number of multiplications and additions for filters with larger support.

Later Schuler and Laine [114, 157] used Simoncelli's construction (originally designed for image coding) for feature detection in mammograms.

Finally, the vision group at NASA [172] proposed a hexagonal image pyramid where successive scales are related by a factor of $\sqrt{7}$. They also noted the similarity of the 7 basis functions of their image pyramid with the receptive fields in the primary visual cortex. However, the fact that each scale change produces a slight rotation of the axes may make this construction quite impractical.

6.4 Biorthogonal hexagonal wavelets

Albert Cohen [49], on the other hand, stresses that it is difficult to obtain orthogonal wavelet bases with compact support, without the help of the tensor product. The problem attributed to the construction of Meyer and Jaffard is that it has no compact support.

Based on these arguments he moves to the biorthogonal case, and derives a set of conditions for constructing biorthogonal wavelet bases, where both the analysis and reconstruction filters are compactly supported and show hexagonal symmetry. Much in the same way as Meyer, he exemplifies with the spline case, and shows how to obtain arbitrarily regular biorthogonal wavelets from these.

Finally he shows how such hexagonally symmetric bases can be used to resolve the orientation ambiguities that result from the wavelet packets obtained from the classical tensor product construction. In fact, he shows that, unlike the square grid case, by iterating the multiscale decomposition also in the high pass regions, the resulting wavelet-packets actually increase the angular resolution.

6.5 Wavelets on irregular grids

Finally Kovacevic, Vetterli and Daubechies [56, 112], propose a biorthogonal wavelet basis construction for irregular grids based on the "lifting scheme", as well as a specialization of such wavelets for the particular hexagonal case.




Figure A.11: Angular resolution that can be obtained by iterating the hexagonal multiscale decomposition also in high-pass regions. The shaded region indicates the frequency region where one of the subbands is concentrated. (Reproduced from [162, page 176, figure 4.17] ; a similar sketch appeared in [49]).

7 Miscellaneous Operations and Applications

7.1 Computer Graphics primitives

In [174] the author extends classical algorithms for discretizing (by means of the nearest neighbor criterion) graphics primitives such as lines, circles and general curves, so that they also work on a hexagonal grid. He shows that the resulting algorithm has the same complexity as its square-grid counterpart, and that the resulting digitized curves more closely reflect the continuity of the original curves.

7.2 Pattern recognition

Several authors found hexagonal algorithms to be useful for several pattern recognition tasks. Here we cite some examples.

In [44, 46] the authors compare edge detection and relaxation methods for hexagonal and square grids. They concluded that the hexagonal procedure is usually more effective and less dependent on the threshold.

In [74] a hexagonal FFT is used to collect features for a later classification of fingerprint images, whereas in [114] hexagonal QMFs are used as an input step for a neural network classifier of handwritten characters. In both cases, the higher orientation selectivity of the hexagonal grid are crucial to obtain a more effective method than in the square case.

7.3 Operating on Hexagonal Grids

The fact that the basis of the hexagonal grid does not consist of orthogonal vectors make some geometric operations more difficult to formulate. In order to simplify the expression of geometric transformations in the hexagonal grid, it has been proposed to use complex arithmetic [28, 166] or a constrained three-dimensional coordinate system [49, 95], that better exploits the axes of symmetry of the hexagonal grid. Another possibility is to use homogeneous coordinates for a special projective basis of the plane, which also exploits these symmetries.

Among the different possibilities mentioned above, the most practical seems to be the constrained three-dimensional coordinate system proposed by Her. Let Γ be a hexagonal grid with basis $\{e_1, e_2\}$. Then a point $x \in \mathbb{R}^2$ can be expressed in the grid coordinates as:

$$x = x_1 e_1 + x_2 e_2 \tag{A.45}$$

Now let us introduce a second set of vectors $\{\tilde{e}_1, \tilde{e}_2, \tilde{e}_3\}$, which we shall call the Her basis, and must satisfy the following conditions:

$$\|\tilde{e}_1\| = \|\tilde{e}_2\| = \|\tilde{e}_3\| \quad (\text{A.46})$$

$$\tilde{e}_i \cdot \tilde{e}_{(i \bmod 3)+1} = cte \quad (\text{A.47})$$

As a consequence all 3 vectors form an angle of 120 degrees and consequently $\tilde{e}_1 + \tilde{e}_2 + \tilde{e}_3 = 0$.

Since Γ is a regular hexagonal grid, we can assume without loss of generality that e_1 and e_2 form an angle of 60 degrees. In such a case we can choose the Her basis such that:

$$\tilde{e}_i \cdot e_j = (1 - \delta_{ij}) \quad \text{for } i, j \in \{1, 2\} \quad (\text{A.48})$$

With these definitions, the Her coordinates of a point are defined by

$$\tilde{x}_i = \langle x, \tilde{e}_i \rangle \quad \text{for } i \in 1, 2, 3. \quad (\text{A.49})$$

and according to the previous relations they are related to the regular coordinates by:

$$\tilde{x}_1 = x_2 \quad (\text{A.50})$$

$$\tilde{x}_2 = x_1 \quad (\text{A.51})$$

$$\tilde{x}_3 = -x_1 - x_2. \quad (\text{A.52})$$

Similarly, for Γ^* we can define a Her coordinate system $\{\tilde{e}_1^*, \tilde{e}_2^*, \tilde{e}_3^*\}$, which will satisfy the same conditions:

$$\|\tilde{e}_1^*\| = \|\tilde{e}_2^*\| = \|\tilde{e}_3^*\| \quad (\text{A.53})$$

$$\tilde{e}_i^* \cdot \tilde{e}_{(i \bmod 3)+1}^* = cte \quad (\text{A.54})$$

However the vectors forming the basis of Γ^* will form an angle of 120 degrees, and the additional condition we can impose to determine the Her basis is:

$$\tilde{e}_i^* \cdot e_j^* = (-1)^i (1 - \delta_{ij}) \quad \text{for } i, j \in \{1, 2\} \quad (\text{A.55})$$

With these definitions, the Her coordinates of a point $\omega = \omega_1 \tilde{e}_1^* + \omega_2 \tilde{e}_2^*$ are defined by

$$\tilde{\omega}_i = \langle x, \tilde{e}_i^* \rangle \quad \text{for } i \in 1, 2, 3. \quad (\text{A.56})$$

and according to the previous relations they are related to the regular coordinates by:

$$\tilde{\omega}_1 = -\omega_2 \quad (\text{A.57})$$

$$\tilde{\omega}_2 = \omega_1 \quad (\text{A.58})$$

$$\tilde{\omega}_3 = \omega_2 - \omega_1 \quad (\text{A.59})$$

7.4 Hardware support for parallel hexagonal image processing

Most parallel processing architectures (most notably hypercubes) are well suited for efficiently performing operations which follow the connectivity pattern of square grids, where each pixel has either 4 or 8 neighbors.

Since Golay [82], however, different hardware architectures have been proposed for efficiently parallelizing algorithms involving operators on 6-neighborhoods. The recent paper by Staunton [168] gives an extensive review of the different possibilities that have been explored.

8 Summary and Discussion

In this appendix we reviewed all the possible ways to sample the image plane, and singled out the square and hexagonal grids as the only reasonable choices. In addition the hexagonal grid was shown to possess several advantages over the more traditional square grid. These advantages were shown first from a theoretical point of view, concluding that the hexagonal grid is the most compact and efficient disposition of the image samples, that it can handle much more elegantly the concept of connectivity, which presents several problems in the square grid.

Secondly, we presented several algorithms and applications from different domains of image processing which show the superiority of the hexagonal grid in certain situations.

Unfortunately, however, two facts have long prevented the extensive use of hexagonal grids and algorithms, namely: the fact that display, acquisition and processing hardware is most often better suited for square grids; and the increased complexity of the algorithms under certain circumstances.

With respect to the second point, it may have become apparent from this review that the increased complexity is only apparent, and that in certain situations hexagonal algorithms are actually not only more simple, but also more effective and efficient.

With respect to the first point we implemented a novel efficient and exact algorithm to convert images between hexagonal and square grids. The proposed algorithm differs from widely used previous approaches in the fact that it preserves the geometry of both regular square and hexagonal grids. Hence, the method could be used to change (without losing information) from one representation to the other, whenever needed to apply an algorithm in its most convenient representation.

Appendix

A Algorithm for separable conversion between Square and Hexagonal Grids

In section 5.4 we briefly described the procedure for resampling on a hexagonal grid Γ^{hex} an image f^{sq} that was sampled on a square grid Γ^{sq} . Recall that the square grid Γ^{sq} has basis $\{\mathbf{u}_1, \mathbf{u}_2\}$ and the target hexagonal grid Γ^{hex} is generated by the basis $\{\mathbf{e}_1, \mathbf{e}_2\}$, which we choose in such a way that $\mathbf{e}_1 = \lambda \mathbf{u}_1$ the first components of the basis are parallel, with proportionality constant λ . In practice:

$$\mathbf{u}_1 = T \begin{pmatrix} 1 \\ 0 \end{pmatrix} \qquad \mathbf{e}_1 = \lambda T \begin{pmatrix} 1 \\ 0 \end{pmatrix} \qquad (\text{A.60})$$

$$\mathbf{u}_2 = T \begin{pmatrix} 0 \\ 1 \end{pmatrix} \qquad \mathbf{e}_2 = \lambda T \begin{pmatrix} 1/2 \\ \sqrt{3}/2 \end{pmatrix} \qquad (\text{A.61})$$

and the proportionality constant $\lambda = \frac{4}{3+\sqrt{3}}$ follows from a trigonometric calculation from the constraint that the hexagon D_{hex}^* be the smallest one containing the square D_{sq}^* . This way Shannon's reconstruction formula will still be valid for the resampled image Γ^{hex} .²

The square to hexagonal conversion procedure consists of two steps:

First, we consider an intermediate grid Γ_1 generated by the basis vectors $\{\mathbf{u}_1, \lambda_1 \mathbf{u}_2\}$. This is just an oversampling in the vertical direction (parallel to \mathbf{u}_2), so that Γ_1 is rectangular. The factor $\lambda_1 = \frac{1}{\|\mathbf{u}_2\|} \langle \mathbf{e}_2, \frac{\mathbf{u}_2}{\|\mathbf{u}_2\|} \rangle$ ensures that the oversampling rate is exactly the one needed for the horizontal lines of Γ_1 to coincide with those of Γ^{hex} . Since we are only changing one of the vectors of the basis, the resampling

$$f_1 = \Gamma_1 \cdot \left(\frac{1}{|D_{\text{sq}}^*|} \overline{\mathcal{F}}(\mathbb{1}_{D_{\text{sq}}^*}) * f^{\text{sq}} \right) \qquad (\text{A.62})$$

reduces to a one-dimensional convolution along each column of $f^{\text{sq}} = \Gamma^{\text{sq}} \cdot f$. Furthermore, since we are just oversampling ($\lambda_1 < 1$), f can still be reconstructed from f_1 by convolution with the same sinc filter.

In the second step we resample f_1 at the target hexagonal grid

$$f^{\text{hex}} = \Gamma^{\text{hex}} \cdot \left(\frac{1}{|D_{\text{sq}}^*|} \overline{\mathcal{F}}(\mathbb{1}_{D_{\text{sq}}^*}) * f_1 \right). \qquad (\text{A.63})$$

² From the arguments in chapter 1, we could slightly relax this condition (and oversample by a factor closer to 1) by imposing that D_{sq}^* be a tile for the dual of Γ^{hex} , and still obtain perfect reconstruction with Shannon's interpolation formula. However, if the hexagonal image is later directly visualized without taking into account the real spectral contents of f , the corners of D_{sq}^* may visually produce some aliasing.

which consists once again only of 1-D convolutions along each of the horizontal lines of Γ_1 . In even lines the resampling is just a zoom, whereas in odd lines the resampling is a zoom plus a translation of half a pixel.

The above procedure is well known. We just rewrote it with our notation, because a similar (possibly less known) procedure can be used to convert to a hexagonal grid Γ^{hex} , an image f^{hex} that was sampled on a square grid Γ^{sq} . As before, the original hexagonal grid Γ^{hex} has basis $\{e_1, e_2\}$ and the target square grid Γ^{sq} is generated by the basis $\{u_1, u_2\}$, which we choose in such a way that $u_1 = \lambda e_1$ the first components of the basis are parallel, with proportionality constant λ . In practice:

$$u_1 = \lambda T \begin{pmatrix} 1 \\ 0 \end{pmatrix} \quad e_1 = T \begin{pmatrix} 1 \\ 0 \end{pmatrix} \quad (\text{A.64})$$

$$u_2 = \lambda T \begin{pmatrix} 0 \\ 1 \end{pmatrix} \quad e_2 = T \begin{pmatrix} 1/2 \\ \sqrt{3}/2 \end{pmatrix} \quad (\text{A.65})$$

and the proportionality constant $\lambda = \frac{3}{4}$ follows from a trigonometric calculation from the constraint that the square Voronoi cell D_{sq}^* of the dual of Γ^{sq} be the smallest one containing the hexagonal Voronoi cell D_{hex}^* of the dual of Γ^{hex} . This way Shannon's reconstruction formula will still be valid for the resampled image Γ^{sq} .

Since, in the hexagonal case, sinc-convolution is 3-separable, in addition to the three basis vectors, it will be convenient to consider a third vector $e_3 = e_2 - e_1$.

The hexagonal to square conversion procedure consists of three steps: First, we consider an intermediate grid Γ_1 generated by the basis vectors $\{e_3, \lambda_1 e_2\}$. This can be seen as an oversampling along the direction e_2 by a factor $\lambda_1 = \|\mathbf{u}_2\| / \langle e_2, \frac{\mathbf{u}_2}{\|\mathbf{u}_2\|} \rangle = \frac{\sqrt{3}}{2}$. This choice of λ_1 will become clear in the second step. Then we compute the resampling

$$f_1 = \Gamma_1 \cdot \left(\frac{1}{|D_{\text{hex}}^*|} \overline{\mathcal{F}}(\mathbb{1}_{D_{\text{hex}}^*}) * f^{\text{hex}} \right) \quad (\text{A.66})$$

which only involves a one-dimensional convolution along each line of $f^{\text{hex}} = \Gamma^{\text{hex}} \cdot f$ parallel to e_2 . Furthermore, since we are just oversampling ($\lambda_1 < 1$), f can still be reconstructed from f_1 by convolution with the same sinc filter.

For the second step we consider an intermediate grid Γ_2 generated by the basis vectors $\{\lambda_1 e_3, \lambda_1 e_2\}$. This can be seen as an oversampling along the direction e_3 by a factor λ_1 . The combination of the first and second steps transformed Γ^{hex} into $\Gamma_2 = \lambda_1 \Gamma^{\text{hex}}$, *i.e.* the horizontal lines of the grid became closer together by a factor λ_1 . The choice of λ_1 ensures that these horizontal lines of Γ_2 coincide with those of the target grid Γ^{sq} . Then, as usual, we compute the resampling

$$f_2 = \Gamma_2 \cdot \left(\frac{1}{|D_{\text{hex}}^*|} \overline{\mathcal{F}}(\mathbb{1}_{D_{\text{hex}}^*}) * f_1 \right) \quad (\text{A.67})$$

which only involves a one-dimensional convolution along each line of f_1 parallel to e_3 . Furthermore, since we are just oversampling ($\lambda_1 < 1$), f can still be reconstructed from f_2 by convolution with the same sinc filter.

In the third step we resample on the grid Γ^{sq}

$$f^{\text{sq}} = \Gamma^{\text{sq}} \cdot \left(\frac{1}{|D_{\text{hex}}^*|} \overline{\mathcal{F}}(\mathbb{1}_{D_{\text{hex}}^*}) * f_2 \right) \quad (\text{A.68})$$

using one-dimensional convolutions along the direction e_1 . Even lines only require a zoom by a factor $\lambda_3 = \frac{\sqrt{3}}{2}$, whereas odd lines require this zoom plus a half-integer shift.

Bibliography

- [1] *Partial Differential Equations*. Chelsea, New York, 1986. see page 275.
- [2] Akram Aldroubi and Hans Feichtinger. Exact iterative reconstruction algorithm for multivariate irregularly sampled functions in spline-like spaces: the L^p theory. *Proceedings of the American Mathematical Society*, 126(9):2677–2686, September 1998.
- [3] Akram Aldroubi and Karlheinz Gröchenig. Beurling-Landau type theorems for non-uniform sampling in shift invariant spline spaces. *The Journal of Fourier Analysis and Applications*, 6(1):93–103, 2000.
- [4] Akram Aldroubi and Karlheinz Gröchenig. Non-uniform sampling and reconstruction in shift-invariant spaces. *SIAM Rev.*, 2001. URL <http://www.math.vanderbilt.edu/~aldroubi/preprints.html#nussis>.
- [5] A. Almansa, J.M. Morel, and B. Rougé. Hexagonal sampling in image processing: A review. Technical report, CMLA, ENS Cachan, August 1999.
- [6] Andrés Almansa. Hexagonal sampling in image processing : A review. Technical report, Centre de Mathématiques et Leurs Applications (CMLA), École Normale Supérieure de Cachan, 94235 Cachan cedex, France, August 1999. Mémoire de stage de DEA.
- [7] Andrés Almansa, Frédéric Cao, Yann Gousseau, and Bernard Rougé. Interpolation of digital elevation models using AMLE and related methods. *IEEE Trans. Geoscience and Remote Sensing*, 40(2):314–325, February 2002.
- [8] Andrés Almansa, Frédéric Cao, Yann Gousseau, and Bernard Rougé. Interpolation of digital elevation models via partial differential equations. *IEEE Trans. on Geoscience and Remote Sensing*, 40(2):314–325, February 2002. submitted Nov 29, 2000; revised Aug 28, 2001.

- [9] Andrés Almansa, Frédéric Cao, and Yann Gousseau and Jean-Michel Morel. Filtres récursifs, déocclusion, interpolation morphologique. Comparaison de méthodes, application aux trous de MNT et bord d'image, réalisation de logiciels relatifs à l'étude. Technical report, CMLA, May 2000. Rapport Technique du CMLA pour le CNES. Marché N° 762/99/CNES/7803. Action de R et T OT102*19.
- [10] Andrés Almansa, Frédéric Cao, and Bernard Rougé. Image interpolation and application to digital elevation models. In *European Workshop Frontiers of Mathematical Morphology*. Université de Strasbourg, April 2000.
- [11] Andrés Almansa and Laurent Cohen. Fingerprint image matching by minimization of thin-plate energy using a two-step algorithm with auxiliary variables. In *Proc. Fifth IEEE Workshop on Applications of Computer Vision (WACV 2000)*, pages 35–40, Palm Springs, CA, December 2000. IEEE, Computer Society Press.
- [12] Andrés Almansa, Agnès Desolneux, Lionel Moisan, and Jean-Michel Morel. Vanishing point detection using helmholtz principle. In *The Mathematical, Computational and Biological Study of Vision*, pages 2–3. Mathematisches Forschungsinstitut Oberwolfach, November 2001. URL http://www.mfo.de/Meetings/Documents/2001/45/Report_49_01.ps.
- [13] Andrés Almansa, Agnès Desolneux, and Sébastien Vamech. Vanishing points are meaningful Gestalts. Technical Report CMLA-2001-24, CMLA, ENS Cachan, 2000.
- [14] Andrés Almansa, Agnès Desolneux, and Sébastien Vamech. Vanishing points are meaningful Gestalts. Technical report, CMLA, ENS Cachan, 2001. URL <http://www.cmla.ens-cachan.fr/Cmla/Publications/2001/CMLA2001-24.pdf>.
- [15] Andrés Almansa, Agnès Desolneux, and Sébastien Vamech. Vanishing point detection without any a priori information. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 2002. special issue on "Perceptual Organization in Computer Vision", to appear.
- [16] Andrés Almansa, Sylvain Durand, and Jean-Michel Morel. Analyse de la résolution et de la stabilité des différents modes d'acquisition SPOT. Technical report, Centre de Mathématiques et Leurs Applications (CMLA), École Nor-

- male Supérieure de Cachan, 94235 Cachan cedex, France, February 2000. Internal Technical Report for CNES.
- [17] Andrés Almansa, Sylvain Durand, and Bernard Rougé. Measuring and improving image resolution by adaptation of the reciprocal cell. 2002. Submitted. Preprint available at <http://www.cmla.ens-cachan.fr/Cmla/Publications/2002/CMLA2002-12.pdf>.
- [18] Andrés Almansa, Stéphane Jaffard, and Jean-Michel Morel. Perturbed sampling in satellite images and reconstruction algorithms. In *18th Symposium GRETSI*, September 2001.
- [19] Andrés Almansa and Tony Lindeberg. Fingerprint enhancement by shape adaptation of scale-space operators with automatic scale selection. *IEEE Trans. on Image Processing*, 9(12):2027–2042, December 2000. Preprint available at <http://www.nada.kth.se/cvap/abstracts/cvap226.html>. Extended version available at <http://www.nada.kth.se/cvap/abstracts/cvap216.html>. Submitted Dec 9, 1998. Revised Jun 27, 2000.
- [20] Yali Amit and Donald Geman. A computational model for visual selection. *Neural Computation*, 1999. URL <http://www.math.umass.edu/~geman/Papers/selection.ps.gz>.
- [21] Richard Andrews and Thong Nguyen. Separable versus quincunx wavelet transforms for image compression. In *6th International Conference on Communications Systems*, Singapore, November 1998. <http://www.eecs.utas.edu.au/people/andrews/apcc98.ps.gz>.
- [22] Richard Andrews and Thong Nguyen. Separable versus quincunx wavelets for image coding. In *6th IEEE International Workshop on Signal Processing and Communications Systems*, Melbourne, Australia, November 1998. <http://www.eecs.utas.edu.au/people/andrews/wispcs98.ps.gz>.
- [23] Mathew E. Antone and Seth Teller. Automatic recovery of relative camera rotations for urban scenes. In *Proceedings of CVPR 2000*, pages II:282–289, 2000.
- [24] G. Aronsson. Extensions of functions satisfying Lipschitz conditions. *Arkiv för Matematik*, 6(28), 1967.

- [25] G. Aronsson. On the partial differential equation $u_x^2 u_{xx} + 2u_x u_y u_{xy} + u_y^2 u_{yy}$. *Arkiv för Matematik*, 7(28), 1968.
- [26] C. Ballester, M. Bertalmío, V. Caselles, G. Sapiro, and J. Verdera. Filling-in by joint interpolation of vector fields and gray levels. *IEEE Trans. Image Processing*, August 2001.
- [27] B. Bauer. *Rappresentazione di Vettori in Spazi ad n Dimensioni*. PhD thesis, Università degli Studi di Milano, 1981.
- [28] S.B.M Bell, F.C Holroyd, and D.C. Mason. A digital geometry for hexagonal pixels. *Image and Vision Computing*, 7(3):194–204, August 1989.
- [29] Marcel Berger. *Géométrie*, volume 1. Nathan, 1990. ISBN 209 191 730-3.
- [30] Marcel Berger. *Géométrie*, volume 2. Nathan, 1990. ISBN 209 191 731-1.
- [31] A. Beurling. . In L. Carleson, P. Malliavin, J. Neuberger, and J. Wermer, editors, *The collected works of Arne Beurling*, volume 2, pages 341–365. Birkhauser, Boston, MA, 1989.
- [32] G. Beylkin. On the fast Fourier transform of functions with singularities. *ACHA*, 2:363–381, 1995.
- [33] G. Beylkin. On applications of unequally spaced fast Fourier transform. In *Mathematical Geophysics Summer School*, Stanford, August 1998.
- [34] Lucien M. Biberman. *Perception of Displayed Information*. Optical Physics and Engineering. Plenum Press, New York, London, 1973.
- [35] T. Boggio. Sull'equilibrio delle piastre elastiche incastrate. *Rend. Acc. Lincei*, 10: 197–205, 1901.
- [36] Jean-Michel Bony. *Cours d'analyse. Théorie des distributions et analyse de Fourier*. Les Éditions de l'École Polytechnique, 2001.
- [37] Fred L. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 11(6):567–585, June 1989.
- [38] G. Borgfors. Distance transformations in hexagonal grids. *Pattern Recognition Letters*, (9):97–105, 1989.

- [39] Y. Bossard. *Rosaces frises et pavages*. CEDIC, 1977.
- [40] Ronald N. Bracewell. *Two-Dimensional Imaging*. Signal Processing Series. Prentice Hall, Englewood Cliffs, New Jersey 07632, 1995.
- [41] F. Cao. Absolutely minimizing Lipschitz extension with discontinuous boundary data. *C.R. Acad. Sci. Paris*, 327:563–568, 1998.
- [42] V. Caselles, J.M. Morel, and C. Sbert. An axiomatic approach to image interpolation. *IEEE Trans. Image Processing*, 7(3), March 1998.
- [43] J-M. Chassery and A. Montanvert. *Géométrie Discrète en analyse d'images*. Hermès, 1991.
- [44] T.I. Cho and K.H. Park. Hexagonal edge relaxation. *Electronics Letters*, 28(4): 357–358, February 1992.
- [45] C. Chui and X. Shi. On stability bounds of perturbed multivariate trigonometric systems. *Applied and Computational Harmonic Analysis*, 3:283–287, 1996.
- [46] In Gook Chun, Taeg Il Cho, and Kyu Ho Park. Boundary detection in a hexagonal grid using energy minimization. *Pattern Recognition Letters*, (15):151–159, February 1994.
- [47] C.V. Coffman. On the structure of solutions of $\Delta^2 u = \lambda u$ which satisfy the clamped plate condition on a right angle. *SIAM Journal of Mathematical Analysis*, 13:746–757, 1982.
- [48] A. Cohen, I. Daubechies, and J.-C. Feauveau. Biorthogonal bases of compactly supported wavelets. *Comm. Pure and Applied Mathematics*, 45:485–560, 1992.
- [49] A. Cohen and J.M. Schlenker. Compactly supported bidimensional wavelet bases with hexagonal symmetry. *Constructive Approximation*, (9):209–236, 1993.
- [50] J. H. Conway and N. J. A. Sloane. *Sphere Packings, Lattices, and Groups*. Springer-Verlag, New York, 2nd ed. edition, 1993.
- [51] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley Series in Telecommunications. John Wiley and sons, 1991.
- [52] M.G. Crandall, H. Ishii, and P.L. Lions. User's guide to viscosity solution of second order partial differential equations. *Bulletin of the American Mathematical Society*, 27(1), July 1992.

- [53] N. Cressie. *Statistics for Spatial Data*. Wiley, New York, revised edition, 1993.
- [54] A. Criminisi, I. Reid, and A. Zisserman. Single view metrology. *International Journal of Computer Vision*, 40(2):123–148, November 2000.
- [55] Germund Dahlquist and Åke Björk. *Numerical methods*. Prentice Hall, Inc., Englewood Cliffs, New Jersey, 1974. isbn 0-13-627315-7.
- [56] I. Daubechies, I. Guskov, P. Schröder, and W. Sweldens. Wavelets on irregular point sets. *Phil. Trans. Royal Society of London A*, 1999. Submitted.
- [57] Julie Delon and Bernard Rougé. Le phénomène d’adhérence en stéréoscopie dépend du critère de corrélation. In *Proc. of GRETSI’01*, Toulouse, September 2001.
- [58] A. Desolneux, S. Ladjal, L. Moisan, and J.M. Morel. Dequantizing image orientation. Technical Report 2000-23, CMLA, ENS Cachan, 2000.
- [59] Agnès Desolneux. *Événements significatifs et applications à l’analyse d’images*. PhD thesis, École Normale Supérieure de Cachan, December 2000.
- [60] Agnès Desolneux, Lionel Moisan, and Jean-Michel Morel. Maximal meaningful events and applications to image analysis. Technical report, preprint CMLA N° 2000-22, 2000. URL <http://www.cmla.ens-cachan.fr/Cmla/Publications/2000/Abstract2000-22.htm>
- [61] Agnès Desolneux, Lionel Moisan, and Jean-Michel Morel. Meaningful alignments. *International Journal of Computer Vision*, 40(1):7–23, 2000.
- [62] Agnès Desolneux, Lionel Moisan, and Jean-Michel Morel. Edge detection by Helmholtz principle. *Journal of Mathematical Imaging and Vision*, 14(3):271–284, 2001.
- [63] Agnès Desolneux, Lionel Moisan, and Jean-Michel Morel. Gestalt theory and computer vision. Technical report, preprint CMLA N° 2002-06, 2002.
- [64] Agnès Desolneux, Lionel Moisan, and Jean-Michel Morel. A grouping principle and four applications. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 2002. special issue on “Perceptual Organization in Computer Vision”, to appear.

- [65] I. Capuzzo Dolcetta and P.L. Lions, editors. *Viscosity solutions and applications*, Lecture Notes in Mathematics, Montecatini Terme, Italy, June 1995. Lectures given at the 2nd session of the Centro Internazionale Matematico Estivo (CIME).
- [66] J. Duchon. Interpolation des fonctions de deux variables suivant le principe de la flexion des plaques minces. *R.A.I.R.O. Analyse numérique*, 10:5–12, 1976.
- [67] Sylvain Durand. Evaluation quantitative de la qualité image. Technical report, CMLA, ENS Cachan, 1998.
- [68] J.C. Ehrhardt. Hexagonal fast Fourier transform with rectangular output. *IEEE Transactions on Signal Processing*, 41(3):1469–1472, March 1993.
- [69] C.L. Fales, F.O. Huck, and R.W. Samms. Imaging system design for improved information capacity. *Applied Optics*, 23(6):872–888, March 1984.
- [70] Olivier Faugeras, Quang-Tuan Luong, and T Papadopoulos. *The Geometry of Multiple Images*. The MIT Press, 2001.
- [71] S.J. Favier and R.A. Zalik. On stability of frames and Riesz bases. *Applied and Computational Harmonic Analysis*, 2:160–173, 1995.
- [72] Hans G. Feichtinger, Karlheinz Gröchenig, and Thomas Strohmer. Efficient numerical methods in non-uniform sampling theory. *Numerische Mathematik*, 69: 423–440, 1995.
- [73] G. Fejes Tóth. Über einen geometrischen Satz. *Math. Z.*, pages 78–83, 1940.
- [74] A.P. Fitz and R.J. Green. Fingerprint classification using hexagonal fast Fourier transform. *Pattern Recognition*, 29(10):1587–1597, 1996.
- [75] D.A.F. Florencio and R.W. Schafer. Homotopy and critical morphological sampling. In *Proceedings of the SPIE's 1994 International Symposium on Visual Communication and Image Processing*, 1994. <http://users.ece.gatech.edu:80/floren/PAPERS/vcip94a/abstract.html>.
- [76] Kristin M. Flornes, Yurii Lyubarskii, and Kristian Seip. A direct interpolation method for irregular sampling. *Applied and Computational Harmonic Analysis*, 7: 315–314, 1999.
- [77] Richard Franke. Scattered data interpolation: Test of some methods. *Mathematics of Computation*, 38(157):181–200, January 1982.

- [78] William T. Freeman, Thonis R. Jones, and Egon C. Pasztor. Example-based super-resolution. Technical report, MERL – A Mitsubishi Electric Research Laboratory, <http://www.merl.com>, August 2001.
- [79] P.R. Garabedian. A partial differential equation arising in conformal mapping. *Pacific J. Math.*, 1:485–524, 1951.
- [80] Claude Gasquet and Patrick Witomski. *Analyse de Fourier et Applications*. Masson, 1995.
- [81] I. M. Gel'fand and N. Y. Vilenkin. *Generalized Functions*, volume 4. Academic Press, New York, 1964.
- [82] Marcel J.E. Golay. Hexagonal parallel pattern transformations. *IEEE Transactions on Computers*, C-18(8):733–740, August 1969.
- [83] Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. The Johns Hopkins University Press, third edition edition, 1996. isbn 0801854148.
- [84] Rafael C. Gonzalez and Richard E. Woods. *Digital Image Processing*. Addison Wesley, 1992.
- [85] K. Gröchenig. Reconstruction algorithms in irregular sampling. *Math. Comp.*, 59 (181–1924), 1992.
- [86] Karlheinz Gröchenig. Irregular sampling , Toeplitz matrices, and the approximation of entire functions of the exponential type. *Mathematics of Computation*, 68(226):749–765, 1999.
- [87] Karlheinz Gröchenig and H. Razafinjatovo. On Landau's necessary conditions for sampling and interpolation of bandlimited functions. *J. London Math. Soc. (2)*, 54(3):557–565, 1996.
- [88] Karlheinz Gröchenig and Thomas Strohmer. Numerical and theoretical aspects of non-uniform sampling of band-limited images. In F. Marvasti, editor, *Theory and Practice of Nonuniform Sampling*. Kluwer/Plenum, 2000.
- [89] Christoph Grunau and Guido Sweers. Positivity for perturbations of polyharmonic operators with Dirichlet boundary conditions in two dimensions. *Math. Nachr.*, 179:89–102, 1996.

- [90] B. Grünbaum and G.C. Shephard. *Tilings and Patterns, an introduction*. W.H Freeman and Company, 1989.
- [91] Carfantan H. and Rougé B. Estimation non biaisée de décalages subpixelaire sur les images SPOT. In *GRETSI*, Toulouse, September 2001.
- [92] J. Hadamard. Mémoire sur le problème d'analyse relatif à l'équilibre des plaques élastiques incastées. *Mémoires Présentées par Divers Savants a l'Académie des Sciences*, 33:1–128, 1908.
- [93] R.M. Haralick, X. Zhuang, C. Lim, and J.S.J. Lee. The digital morphological sampling theorem. *IEEE Transactions on Acoustics, Speech and Signal Processing*, 37: 2067–2090, December 1989.
- [94] Håkan Hendenmalm. An Hadamard maximum principle for the bilaplacian on hyperbolic manifolds. In *Journées Équations aux dérivées partielles*, Saint-Jean-de-Monts, June 1999. GDR 1151 (CNRS).
- [95] I. Her. Geometric transformations on the hexagonal grid. *IEEE Trans. on Image Processing*, 4(9):1213–1222, September 1995.
- [96] Innchyn Her and Chi-Tseng Yuan. Resampling on a pseudo-hexagonal grid. *Computer Vision, Graphics, and Image Processing. Graphical Models and Image Processing*, 56(4):336–347, July 1994. ISSN 1049-9652.
- [97] Stéphane Jaffard. *Construction et propriétés des bases d'ondelettes. Remarques sur la contrôlabilité exacte*. PhD thesis, École Polytechnique, May 1989.
- [98] Stéphane Jaffard. Résultats généraux sur l'échantillonnage irrégulier. Personal communication, 2000.
- [99] Stéphane Jaffard. Beyond Besov spaces. Personal Communication, 2002.
- [100] Bernd Jähne. *Digital Image Processing*. Springer-Verlag, 3rd edition, 1995.
- [101] Anil K. Jain. *Fundamentals of Image Processing*. Prentice Hall, 1995.
- [102] R. Jensen. Uniqueness of Lipschitz extension: minimizing the sup norm of the gradient. *Arch. Rat. Mech. Anal.*, (123):51–74, 1993.
- [103] A. G. Journel and C. J. Huijbregts. *Mining Geostatistics*. Academic Press, London, 1978.

- [104] M. I. Kadec. The exact value of the Payley-Wiener constant. *Soviet Math. Doklady*, 5(559–561), 1964.
- [105] Behrooz Kamgar-Parsi and Behzad Kamgar-Parsi. Quantization error in hexagonal sensory configurations. *IEEE Trans. on Image Processing*, 14(6):665–671, June 1992.
- [106] Behrooz Kamgar-Parsi and Behzad Kamgar-Parsi. Quantization error in regular grids: Triangular pixels. *IEEE Trans. on Image Processing*, 7(10):1496–1500, October 1998.
- [107] Gaetano Kanizsa. *Organization in Vision*. Holt, Rinehart & Winston, 1979.
- [108] Gaetano Kanizsa. *Grammatica del Vedere*. Società editrice Il Mulino, Bologna, 1980.
- [109] Gaetano Kanizsa. *La Grammaire du Voir*. Diderot, 1997.
- [110] J.T. Kent and K.V. Mardia. The link between Kriging and thin-plate splines. In F.P. Kelly, editor, *Probability, Statistics and Optimisation*, chapter 24. John Wiley & Sons, Inc., 1994.
- [111] L. Klaas. Optimale Abtastung reeller Funktionen in der euklidischen Ebene. *Archiv für Elektronik und Übertragungstechnik*, 39(1):57–60, 1985.
- [112] J. Kovacevic and W. Sweldens. Wavelet families of increasing order in arbitrary dimensions. *IEEE Trans. on Image Processing*, 1998. Submitted.
- [113] J. Kovacevic, M. Vetterli, and G. Karlson. Design of multidimensional filter banks for non-separable sampling. In *IEEE International Symposium on Circuits and Systems*, pages 2004–2008, 1990.
- [114] A. F. Laine and S. Schuler. Hexagonal wavelet representations for recognizing complex annotations. In *Proceedings of the Conference on Computer Vision and Pattern Recognition*, pages 740–745, Los Alamitos, CA, USA, June 1994. IEEE Computer Society Press. ISBN 0-8186-5825-8.
- [115] H.J. Landau. Necessary density conditions for sampling and interpolation of entire functions. *Acta Math*, 117:37–52, 1967.
- [116] Christophe Latry, Bernard Rougé, and Baillarin S. La chaîne image SPOT5 THR : un exemple d’optimisation globale. In *GRETSI*, Toulouse, September 2001.

- [117] M. Li, M.J. Philips, and D.J. Whitehouse. Extension of two-dimensional sampling theory. *Journal of Physics A: Mathematical and General*, 22:5053–5063, 1989.
- [118] D. Liebowitz, A. Criminisi, and A. Zisserman. Creating architectural models from images. *EuroGraphics*, 18(3):xx–yy, 1999.
- [119] E. Luczak and A. Rosenfeld. Distance on a hexagonal grid. *IEEE Transactions on Computers*, pages 532–533, May 1976.
- [120] Evelyne Lutton, Henri Maître, and Jaime Lopez-Krahe. Contribution to the determination of vanishing points using Hough transform. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(4):430–438, 1994.
- [121] François Malgouyres. *Augmentation de la résolution d'images satellitaires: Théorie variationnelle et applications*. PhD thesis, Université de Paris IX – Dauphine, 1999.
- [122] François Malgouyres and Frederic Guichard. Edge direction preserving image zooming: a mathematical and numerical analysis. *Journal on Numerical Analysis*, 39(1):1–37, 2001.
- [123] S. Masnou. *Filtrage et Desocclusion d'Images par Methodes d'Ensembles de Niveau*. PhD thesis, Universite Paris-IX Dauphine, 1998.
- [124] G. Matheron. La théorie des variables régionalisées, et ses applications. *Les Cahiers du Centre de Morphologie Mathématique de Fontainebleau*, 5, 1971.
- [125] G. Matheron. The intrinsic random functions and their applications. *Advances in Applied Probability*, 5:439–468, 1973.
- [126] G. Matheron. Splines and Kriging; their formal equivalence. In D.F. Merriam, editor, *Down-to-Earth Statistics: Solutions looking for Geological Problems*, pages 77–95. Syracuse University Geology Contributions, 1981.
- [127] Jean Meinguet. Surface spline interpolation: Basic theory and computational aspects. In S.P. Singh et al., editor, *Approximation Theory and Spline Functions*, pages 124–142. D. Reidel Publishing Company, 1984.
- [128] R.M. Mersereau. The processing of hexagonally sampled two-dimensional signals. *IEEE Proceedings*, 67:930–949, 1979.

- [129] R.M. Mersereau and T.C. Speake. The processing of periodically sampled multi-dimensional signals. *IEEE Transactions on Acoustics, Speech and Signal Processing*, 31(1):188–194, February 1983.
- [130] Yves Meyer. *Wavelets and operators*, volume 2. Hermann, 1990.
- [131] Alexey A. Minho and Sergey A. Pokotilo. Nonuniform quantization of the image in infrared remote sensing systems: Ways and devices for its elimination. In *Proceedings of SPIE*, volume 3063, pages 312–320, 1997.
- [132] L. Mitas and H. Mitasova. Spatial interpolation. In Paul A. Longley, Michael F. Goodchild, David J. Maguire, and David W. Rhind, editors, *Geographical Information Systems*, volume 1, chapter 34, pages 481–492. John Wiley & Sons, Inc., second edition edition, 1999.
- [133] Lionel Moisan and Jean-Michel Morel. Les défauts du DVD sous l’oeil de l’équation. *La Recherche*, (342):56–57, May 2001.
- [134] P. Monasse and F. Guichard. Fast computation of contrast invariant image representation. *IEEE Trans. on Image Processing*, 1998. submitted.
- [135] J.M. Morel and S. Ladjal. Analyse de Fourier et théorie de Shannon. Analyse de Fourier et traitement d’images. Journées X-UPS 1998. Centre de Mathématiques – Ecole Polytechnique. France., 1998.
- [136] André Nel. Hexagonal image processing. In *Southern African Conference on Communications and Signal Processing (COMSIG)*, pages 109–113, 1989.
- [137] J. Pipher and G. Verchota. A maximum principle for biharmonic functions in Lipschitz and C^1 domains. *Comment. Math. Helvetici*, 68:385–414, 1993.
- [138] Murphy P.K. and Gallagher N.C. Hexagonal sampling techniques applied to Fourier and Fresnel digital holograms. *Journal of the Optical Society of America*, 72(7):929–37, July 1982.
- [139] William K. Pratt. *Digital Image Processing*. John Wiley and Sons, 1978.
- [140] K. Preston, M.J.B. Duff, S. Levialdi, P.E. Norgren, and J. Toriwaki. Basics of cellular logic with some applications in image processing. *IEEE Proceedings*, 67(5): 826–856, 1979.

- [141] Michael Rauth. *Gridding geophysical potential fields from noisy scattered data*. PhD thesis, University of Vienna, May 1998. URL <http://www.rauth.at/papers/thesis.asp>.
- [142] Sylvie Rocques, Branchère F., and Bernard Rougé. Séparation des décalages induits par l'attitude et le relief entre images d'un couple stéréoscopique. In *GRETSI*, Toulouse, September 2001.
- [143] C.A. Rogers. *Packing and Covering*. Cambridge University Press, 1964.
- [144] J.K. Romberg, Hyeokho Choi, and Richard Baraniuk. Bayesian tree-structured image modeling using wavelet-domain hidden Markov models. *IEEE Trans. on Image Processing*, 10(7):1056–1068, July 2001.
- [145] Azriel Rosenfeld and Avinash C. Kak. *Digital Picture Processing*, volume 1 of *Computer Science and Applied Mathematics*. Academic Press, second edition, 1982.
- [146] C. Rother. A new approach for vanishing point detection in architectural environments. In *British Machine Vision Conference*, pages xx–yy, 2000.
- [147] Bernard Rougé. Théorie de la chaîne image optique et restauration à bruit final fixé. Mémoire en vue de l'obtention de l'habilitation à diriger des recherches. Option: Mathématiques appliquées, May 1997.
- [148] Bernard Rougé. Théorie de la chaîne image optique et restauration à bruit final fixé. Mémoire en vue de l'obtention de l'habilitation à diriger des recherches. Option: Mathématiques appliquées, May 1997.
- [149] Bernard Rougé. Théorie de l'échantillonnage et satellites d'observation de la terre. Analyse de Fourier et traitement d'images. Journées X-UPS 1998. Centre de Mathématiques – Ecole Polytechnique. France., 1998.
- [150] Bernard Rougé. Corrélation fine, interpolation régulière et irrégulière, échantillonnage de Shannon. Internal Report, CNES, Toulouse, France., March 1999.
- [151] Bernard Rougé and Vicent Caselles. Superrésolution couleur a partir d'un canal panchromatique haute résolution et des canaux couleur basse résolution. Personal Communication, 2001.

- [152] L. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992.
- [153] John C. Russ. *The image processing handbook*. CRC, 2nd edition, 1995.
- [154] L. Rudin S. Jensen. Measure: an interactive tool for accurate forensic photo/video grammetry. In *Investigative & Trial Image Processing Conference*, volume 2567 of *Proceedings of SPIE*, San Diego, CA, 1995.
- [155] Luis A. Santaló. Integral geometry and geometric probability. In Gian-Carlo Rota, editor, *Encyclopedia of Mathematics and its Applications*, volume 1. Addison-Wesley, 1976.
- [156] F. Schaffalitzky and A. Zisserman. Planar grouping for automatic detection of vanishing lines and points. *Image and Vision Computing*, 18(9):647–658, June 2000.
- [157] S. Schuler and A. Laine. Hexagonal QMF banks and wavelets. In M. Akai, editor, *Time-Frequency and Wavelet Transforms in Biomedical Engineering*, chapter ? IEEE Press, 1997.
- [158] J. Serra and B. Laÿ. Square to hexagonal lattices conversion. *Signal Processing*, (9):1–13, 1985.
- [159] Jean Serra. *Image Analysis and Mathematical Morphology*, volume 1. Academic Press, 1982.
- [160] J.A. Shufelt. Performance evaluation and analysis of vanishing point detection techniques. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 21(3):282–288, March 1999.
- [161] Eero P. Simoncelli and Edward H. Adelson. Subband image coding with hexagonal quadrature mirror filters. In *Picture Coding Symposium*, Cambridge, MA., March 1990.
- [162] E.P Simoncelli and E.H Adelson. Subband transforms. In John W. Woods, editor, *Subband Coding*, chapter 4, pages 143–192. Kluwer Academic Publishers, 1990. Also available from MIT Media Laboratory as Technical Report #137, <ftp://whitechapel.media.mit.edu/pub/tech-reports/TR-137.ps.Z>.
- [163] Pierre Soille. Spatial distributions from contour lines: An efficient methodology based on distance transformations. *Journal of Visual Communication and Image Representation*, 2(2):138–150, June 1991.

- [164] Pierre Soille. *Morphologie Mathématique : du relief à la dimensionalité – Algorithmes et méthodes*. PhD thesis, Université Catholique de Louvain, Louvain-la-Neuve, Belgium, 1992.
- [165] Pierre Soille. Generalized geodesic distances applied to interpolation and shape description. In Jean Serra and Pierre Soille, editors, *Mathematical Morphology and its Applications to Image Processing*. Kluwer Academic Publishers, 1994.
- [166] R.C. Staunton. The design of hexagonal sampling structures for image digitization and their use with local operators. *Image and Vision Computing*, 7(3):162–166, August 1989.
- [167] R.C. Staunton. An analysis of hexagonal thinning algorithms and skeletal shape representations. *Pattern Recognition*, 29(7):1131–1146, 1996.
- [168] R.C. Staunton. Hexagonal sampling in image processing. *Advances in Imaging and Electron Physics*, 107:231–307, 1999.
- [169] S.T. Barnard. Interpreting perspective images. *Artificial Intell.*, 21:435–462, 1983.
- [170] T. Tuytelaars, M. Proesmans, and L. Van Gool. The cascaded Hough transform. In *International Conf. on Image Processing, ICIP-97*, volume 2, pages 736–739, 1997.
- [171] Michael Unser. Sampling – 50 years after Shannon. *Proceedings of the IEEE*, 88(4):569–587, April 2000.
- [172] A.B. Watson and A.J. Ahumada-Jr. A hexagonal orthogonal-oriented pyramid as a model of image representation in visual cortex. *IEEE Transactions on Biomedical Engineering*, 36(1), January 1989. Also available from <http://vision.arc.nasa.gov/publications/hexagonal.pdf>. Previous version available as NASA TM-100054, 1987, <http://vision.arc.nasa.gov/publications/OrthogonalHexagonal.pdf>.
- [173] Max Wertheimer. Untersuchungen zur Lehre von der Gestalt. II. *Psychologische Forshung*, 4:301–350, 1923.
- [174] C.A. Wühtrich and P. Stucki. An algorithmic comparison of square- and hexagonal-based grids. *CVGIP: Graphical Models and Image Processing*, 53(4):324–339, 1991.

- [175] Robert M. Young. *Introduction to Nonharmonic Fourier Series*. Academic Press, 1980.
- [176] Jaime L. Zapata and Gerhard X. Ritter. Fast Fourier transform for hexagonal aggregates. *Journal of Mathematical Imaging and Vision*, 12:183–197, 2000.

Annexe B

Systemes d'échantillonnage du CNES *

*A preliminary version of this annex was submitted as a technical report to CNES, jointly with Sylvain Durand and Jean-Michel Morel [[16](#)].

Résumé

Cette thèse aborde quelques-uns des problèmes qui surviennent dans la conception d'un système complet de vision par ordinateur : de l'échantillonnage à la détection de structures et leur interprétation. La motivation principale pour traiter ces problèmes a été fournie par le CNES et la conception des satellites d'observation terrestre, ainsi que par les applications de photogrammétrie et vidéo-surveillance chez Cognitech, Inc. pendant les étapes finales de ce travail, mais les techniques développées sont d'une généralité suffisante pour présenter un intérêt dans d'autres systèmes de vision par ordinateur.

Dans une première partie nous abordons une étude comparative des différents systèmes d'échantillonnage d'images sur un réseau régulier, soit carré soit hexagonal, à l'aide d'une mesure de résolution effective, qui permet de déterminer la quantité d'information utile fournie par chaque pixel du réseau, une fois que l'on a séparé les effets du bruit et du repliement spectral. Cette mesure de résolution est utilisée à son tour pour améliorer des techniques de zoom et de restauration basées sur la minimisation de la variation totale. Ensuite l'étude comparative est poursuivie en analysant dans quelle mesure chacun des systèmes permet d'éliminer les perturbations du réseau d'échantillonnage dues aux micro-vibrations du satellite pendant l'acquisition. Après une présentation des limites théoriques du problème, nous comparons les performances des méthodes de reconstruction existantes avec un nouvel algorithme, mieux adapté aux conditions d'échantillonnage du CNES.

Dans une deuxième partie nous nous intéressons à l'interpolation de modèles d'élévation de terrain, dans deux cas particuliers: l'interpolation de lignes de niveau, et l'étude des zones dans lesquelles une méthode de corrélation à partir de paires stéréo ne fournit pas des informations fiables. Nous étudions les liens entre les méthodes classiques utilisées en sciences de la terre tels que Krigeage ou distances géodésiques, et la méthode AMLE, et nous proposons une extension de la théorie axiomatique de l'interpolation qui conduit à cette dernière. Enfin une évaluation expérimentale permet de conclure qu'une nouvelle combinaison du Krigeage avec l'AMLE fournit les meilleures interpolations pour les modèles de terrain.

Enfin nous nous intéressons à la détection d'alignements et de leurs points de fuite dans une image, car ils peuvent être utilisés aussi bien pour la construction de modèles d'élévation urbains, que pour résoudre des problèmes de photogrammétrie et calibration de caméras. Notre approche est basée sur la théorie de la Gestalt, et son implémentation effective récemment proposée par Desolneux-Moisan-Morel à l'aide du principe de Helmholtz. Le résultat est un détecteur de points de fuite sans paramètres, qui n'utilise aucune information a priori sur l'image ou la caméra.

Abstract

This thesis deals with a few of the many issues that arise in the design of a complete computer vision system, from sampling and interpolation, to feature detection and interpretation. The main motivation for addressing these topics was provided by the French centre of space studies (CNES) and the design of earth observation satellites, as well as photogrammetry applications and video-surveillance applications at Cognitech, Inc. during the final stages of this work, but most subjects are treated with sufficient generality to be of interest for other computer vision systems.

In a first part we perform a comparative study of different sampling systems on a regular grid, which can be either square or hexagonal. We do so by means of an effective image resolution measure, which allows to determine the mean amount of useful information contained in each pixel, once the noise and aliasing effects have been discarded. This resolution measure is used to improve the zoom and restoration methods based on total variation minimization. Next the comparative study is continued by analysing to what an extent each sampling system allows to undo the perturbations of the sampling grid due to the satellite micro-vibrations during image acquisition. After a review of the theoretical limits of this reconstruction problem, we compare the performance of available reconstruction methods with a new one, which is better adapted to the sampling conditions of CNES's systems.

In a second part we address the interpolation of digital terrain models in two particular cases: the interpolation of level curves, and of those regions where a stereo-pair correlation method failed to provide reliable height information. We study the links of classical methods used in the geoscience literature, such as Kriging and geodesic distance, with the AMLE method, and we propose an extension of the axiomatic interpolation theory leading to the latter. Finally, an experimental evaluation allows us to conclude that a new combination of Kriging and AMLE provides in general better interpolations for terrain models.

At last, we consider the automatic detection of alignments and their vanishing points in digital images, since they can be used both for constructing elevation models in urban areas, and for solving photogrammetry and camera calibration problems. Our approach is based on Gestalt theory and its computational implementation recently proposed by Desolneux-Moisan-Morel using the Helmholtz principle. The result is a parameterless vanishing point detector, which doesn't require any a priori information about the image or the camera calibration parameters.

