



HAL
open science

Modèles variationnels et bayésiens pour le débruitage d'images : de la variation totale vers les moyennes non-locales

Cécile Louchet

► **To cite this version:**

Cécile Louchet. Modèles variationnels et bayésiens pour le débruitage d'images : de la variation totale vers les moyennes non-locales. Mathématiques [math]. Université René Descartes - Paris V, 2008. Français. NNT: . tel-00371438

HAL Id: tel-00371438

<https://theses.hal.science/tel-00371438v1>

Submitted on 27 Mar 2009

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.

UNIVERSITÉ PARIS DESCARTES
UFR Mathématiques Informatique
École Doctorale Mathématiques Paris-Centre

THÈSE

pour obtenir le grade de

Docteur de l'Université Paris Descartes
Spécialité : Mathématiques Appliquées

présentée par

Cécile LOUCHET

**Modèles variationnels et bayésiens
pour le débruitage d'images :
de la variation totale vers les moyennes non-locales**

**Variational and Bayesian models for image denoising :
from total variation towards non-local means**

Soutenue le 10 décembre 2008 devant le jury composé de :

Laure BLANC-FÉRAUD	INRIA Sophia	Examineur
Vicent CASELLES	Université Pompeu Fabra, Espagne	Examineur
Antonin CHAMBOLLE	École Polytechnique	Rapporteur
Charles KERVRANN	IRISA/INRIA Rennes	Examineur
François MALGOUYRES	Université Paris 13	Examineur
Yves MEYER	ENS Cachan	Examineur
Lionel MOISAN	Université Paris Descartes	Directeur de thèse
Mila NIKOLOVA	ENS Cachan	Rapporteur

Remerciements

En tout premier lieu je voudrais remercier très chaleureusement Lionel Moisan pour ces quelques années passées sous sa direction. Je l'ai toujours admiré, pour ses compétences scientifiques d'abord, mais aussi pour son éclectisme, qui rend sa compagnie très agréable, et ses qualités humaines remarquables dont j'ai vraiment profité! Tout en me laissant une grande liberté dans mon travail, il m'a permis, par sa patience, ses encouragements et son écoute, de traverser sans doute plus facilement les mauvaises passes que tout doctorant est amené à vivre. Pour tout cela, je le remercie du fond du cœur, ainsi que pour toutes les marques de confiance qu'il m'a accordées et la relation amicale que nous avons construite.

Je voudrais remercier les membres du jury. Tout d'abord, Antonin Chambolle et Mila Nikolaeva m'ont fait grand honneur en acceptant de rapporter ma thèse. Ils ont non seulement passé de longues heures, que j'imagine lourdes en maux de crâne, à corriger le manuscrit, mais surtout m'ont proposé de nouvelles références, de nouvelles preuves plus élégantes, avec des perspectives très enthousiasmantes. Chacun d'eux, à sa manière, a su m'insuffler une bonne dose de dynamisme dans mon travail!

Laure Blanc-Féraud a accepté de présider mon jury; qu'elle en soit très sincèrement remerciée. Sa bonne humeur a donné le ton à une soutenance chaleureuse et agréable, ainsi qu'à une discussion scientifique très riche.

Je voudrais remercier Yves Meyer d'avoir accepté de participer au jury. Je me souviens d'un cours de théorie de la mesure où il avait réussi à me faire pleurer d'émotion... Je le remercie chaleureusement pour toute la poésie qu'il a su nous enseigner à travers les mathématiques, et pour sa constante bienveillance.

Je suis très reconnaissante à Charles Kervrann, François Malgouyres ainsi qu'à Vicent Caselles d'avoir accepté de faire partie de mon jury. Je les remercie pour tout l'intérêt qu'ils ont porté à ce travail, et pour leurs remarques constructives, pendant et même après la soutenance.

Je voudrais aussi remercier Gersende Fort pour sa discrète mais importante contribution dans la lecture du manuscrit; ses remarques constructives sont autant de pistes à explorer.

Je voudrais exprimer ma reconnaissance aux quelques chercheurs confirmés qui m'ont aidée dans mes recherches, en m'aiguillant amicalement et dans le service dévoué à la science. Je pense à Servane Gey, Simon Masnou, Annie Raoult, Eva Löcherbach, Fabienne Comte, Gilles Pagès, qui ont pu me tirer quelques épines techniques du pied. Je pense aussi à Erwan Le Penneç, Charles Kervrann, Frédéric Richard et Jean-Michel Morel pour des discussions très riches.

Je voudrais remercier tout spécialement les personnes qui, directement ou indirectement, m'ont toujours donné confiance et encouragée dans mon travail, et qui ont su ainsi faire un travail psychologique intense! Je pense à Gabriel Peyré, Maitine Bergounioux, Claire Lacour, Pierre Calka, Avner Bar-Hen, George Stamon, Julie Delon, et puis tout particulièrement Jean-François Aujol qui m'a beaucoup soutenue pendant ces quelques années.

Les personnes qui ont partagé ma vie quotidienne au labo m'ont énormément apporté tout au long de la thèse. Bernard Ycart puis Christine Graffigne ont été pour moi des directeurs de laboratoire patients et très disponibles; je les remercie chaleureusement pour m'avoir offert de très bonnes conditions matérielles pour faire ma thèse, ainsi que d'avoir su instaurer des relations de confiance entre doctorants et permanents. Merci aussi à Christophe Castellani et à

Pascal Tournois qui m'ont bien aidée dans les aspects administratifs rébarbatifs de la thèse. Mais surtout, je ne peux être que profondément reconnaissante envers toute la clique des thésards du labo : Amandine, Béatrice, Olivier, Claire L., Gwendoline, Claire J., Bénédicte, Javiera, Sylvain, qui ont su m'accueillir au MAP5 très amicalement, ainsi que Mohamed, Gaëlle, Neus, Makrem, Solange, Georges, Alice, Julien, Maël, Nathalie, Mahendra, Maxime, Baptiste, Sandra, Benjamin, Emeline, Arno, avec une spéciale dédicace pour les 5 ou 8 derniers qui ont supporté de plein fouet ma petite forme de fin de thèse. Merci à Claire L. pour son amitié et son bon-sens à toute épreuve qui m'ont toujours remis le moral au top ! Je n'oublie pas non plus Jean-Pascal, Julie, Benjamin, Christophe, Pascal, Béranger, Sébastien, Julien, Frédéric, avec qui j'ai partagé du temps au CMLA (Cachan) en début de thèse. Au CMLA comme au MAP5, le groupe des thésards est bien soudé, et les soirées entre thésards, le groupe de travail, les cafés au bar, tout cela fait que le travail est nettement plus sympa !

Grand merci à ceux qui ont accepté de "pauser" pour ma thèse : David, François, Anne L., Amélie, Anne P., Martin ont accepté de voir leurs œuvres photographiques ou leur visage dans mon manuscrit. Grâce à eux, le manuscrit est un peu plus original ! Merci aussi à ceux qui ont pris le temps de lire la thèse, ou même quelques pages, ils se reconnaîtront.

Je veux tout particulièrement remercier mes proches qui m'ont spécialement "coachée" pendant la rédaction de thèse (en plus de Lionel bien sûr) : David, Benoît, Marie, Béatrice, Jeanne, André et Marie-Christine, mes parents (merci Papa pour les approvisionnements indispensables en chocolat noir !), en plus des thésards, ainsi que Piotr, N.B. et P.C., les plus fidèles. Pensée spéciale à la mairie de Paris qui m'a permis de profiter à vélib' de la tiédeur des nuits d'été parisiennes, quelques minutes par jour ! Et plus sérieusement, je tiens à remercier mes proches qui, dans leur affection, sont restés discrets pendant cette période, et m'ont permis de me concentrer au maximum.

Enfin, un merci tout spécial à ma famille que j'aime, et qui m'a toujours soutenue. En particulier, ma sœur dont j'ai toujours été si proche, André et Marie-Christine qui m'ont "cocoonnée" comme leur propre fille, mes grands-parents toujours prompts à me recevoir pour le repas, même prévenus avec deux minutes d'avance, et mes parents, qui m'ont bien soutenue, l'un avec franchise et scepticisme, l'autre toujours fière de moi, quoi que je fasse... Je n'oublie pas mes amis qui m'ont été d'un réel soutien, et l'humanité tout entière, à qui cette thèse est dédiée !

Table des matières

Introduction	3
I Total Variation denoising using posterior expectation	15
1 Motivation and definition of the TV-LSE denoiser	17
1.1 Total Variation minimization	17
1.2 TV minimization and Maximum a Posteriori	20
1.2.1 A statistical framework for TV minimization	20
1.2.2 Loss functions	22
1.2.3 Distortions of MAP estimation in image denoising	24
Distortion due to high dimension	24
Distortions in simple cases of MAP estimation	25
Staircasing for TV-denoising	27
Temperature problem	31
1.3 Definition of TV-LSE	32
1.3.1 Constrained TV-LSE	33
1.3.2 TV-LSE as a diffusion	34
1.3.3 Link to Wiener filter	34
2 MCMC algorithm for TV-LSE denoising	37
2.1 MCMC and Hastings-Metropolis algorithms	38
2.2 A Hastings-Metropolis algorithm for TV-LSE	41
2.2.1 Subsampled chain	42
2.2.2 Consistency of Algorithm 2	43
2.3 Convergence control	45
2.3.1 Central limit theorem	46
2.3.2 Bias-variance decomposition	48
2.3.3 Introduction of another chain, and stopping criterion	53
2.4 Burn-in selection	54
2.5 Optimal scaling	57
2.6 Numerical variants – why we won't run several chains	61
2.6.1 Several independent MCMCs	61
2.6.2 Interacting MCMCs – Sampling Importance Resampling	63

3	Theoretical properties	67
3.1	Basic properties	68
3.1.1	Average preservation	68
3.1.2	Invariances	69
3.1.3	Regularity	71
3.1.4	Explicit solution for a 2-pixel clique	73
3.2	Asymptotic behavior of TV-LSE	76
3.3	No staircasing in TV-LSE denoised images	81
3.4	Brief analysis of the assumptions on discrete TV	85
3.5	Global analysis of LSE denoisers built on a more general prior	87
4	Numerical experiments	99
4.1	Discussion on the stopping criterion	100
4.2	Influence of parameters σ , λ and β	105
4.3	Denoising normalization and comparisons to other methods	120
4.4	Experiments on synthetic signals	136
4.5	Staircasing and isolated pixels artefacts	141
4.6	Other examples, and natural noise	145
II	Locality and prior	151
5	Locality in Total Variation denoising	153
5.1	Introduction	153
5.2	Total variation minimization is global	154
5.3	Lagrangian local TV-MAP	156
5.3.1	Construction	156
5.3.2	Algorithm	157
5.3.3	Stability properties	165
5.3.4	PDE associated to local TV-MAP	168
5.3.5	Local versus global TV-MAP	173
5.3.6	Bias-variance trade-off for window optimizing	176
5.3.7	Adaptive locality as a function of the gradient norm	183
	Introduction	183
	Estimation of the optimal window by the local scale	183
	Definition of a locally optimal window	184
	Optimal locality and gradient norm	184
	Adaptive denoising	186
5.4	Constrained local TV-MAP	190
5.4.1	Algorithm	191
5.4.2	Minimum size of windows?	192
5.4.3	Experiments	192
5.5	Towards a quadratic risk as in the Non-Local means	198

6	Total Variation prior model, and the Non-Local means	201
6.1	Introduction	202
6.2	Prior in Total Variation denoising	202
6.2.1	Gradient norm is a good measure of small patch distribution	202
6.2.2	Optimizing the total variation prior	206
6.3	Bias-variance trade-off in the prior complexity	211
6.3.1	Bayesian framework for NL-means	211
6.3.2	Bias and variance for TV-denoising and for NL-means	212
6.3.3	Argument for the existence of a non-trivial optimum	213
	Conclusion	216
	Bibliographie	221

Introduction

Dans cette thèse, nous nous intéressons à la réduction du bruit dans les images. Le bruit n'est pas qu'un phénomène acoustique, on peut aussi en trouver dans le domaine visuel. Par exemple quand l'antenne d'un poste de télévision est mal orientée, l'image à l'écran est de très mauvaise qualité et contient des parasites, très désagréables à l'oeil, que l'on appelle du bruit. Le bruit peut donc être naturellement présent dans la vidéo, mais aussi dans les images fixes, prises par exemple avec un appareil photographique. De fait une image peut contenir des dégradations de toutes sortes par rapport à la réalité d'une scène qu'elle est censée représenter. Citons quelques-unes de ces dégradations.

Tout d'abord, une image peut être floue par mise au point, soit parce que le réglage de l'appareil est mal fait, soit parce que la scène contient des objets à des profondeurs différentes, ce qui rend impossible la mise au point. Une autre source possible de flou provient du bougé, quand le temps d'obturation est long et que la scène ou la caméra est mobile.

Une autre source de détérioration est liée aux défauts de réalisation de l'optique, qui est toujours le résultat d'un compromis difficile. Certaines aberrations (comma, astigmatisme) affectent la netteté de l'image ; d'autres aberrations concernent la géométrie ou la radiométrie, comme le vignetage (réduction de l'ouverture relative en périphérie de l'image) ou les réflexions internes parasites. De plus, la miniaturisation des matrices de capteurs limite la qualité des images à cause du phénomène de diffraction.

Si le réglage de la lumière est mal fait ou si la scène est trop contrastée par rapport à la dynamique du capteur (ou à la dynamique imposée par le codage numérique choisi), la réponse du capteur peut devenir fortement non-linéaire, et on peut avoir sous-exposition ou sur-exposition. Alors il peut arriver que des niveaux de gris saturent dans le noir ou le blanc ; ainsi, dans les zones saturées, l'information présente dans la scène est perdue dans la photographie.

Le passage d'une scène définie dans un espace continu à une grille discrète peut générer des déformations importantes, dues à la perte d'information, comme par exemple des effets de moirés (*aliasing*) ou d'écho (*ringing*). Enfin, des déformations dues à la compression des images numériques, par exemple les blocs JPEG, altèrent aussi la qualité de l'image quand le taux de compression est poussé.

Dans cette thèse, nous nous intéressons au bruit radiométrique. Une source importante de bruit trouve son origine dans le capteur. On connaît bien le « grain photographique » présent dans les capteurs argentiques traditionnels, surtout avec les sensibilités élevées ; les capteurs électroniques ne sont pas exempts de ce genre de défaut. Les capteurs électroniques ont une précision limitée dans le comptage de photons, et introduisent ainsi un bruit de quantification, particulièrement présent lorsque de faibles éclaircissements sont en jeu. Un trop grand contraste de

la scène (contre-jour, par exemple) peut entraîner des difficultés dans le transfert des charges électroniques et un éblouissement du capteur, donnant lieu à un bruit. Scanner une image analogique ajoute du bruit; rehausser le contraste d'une photo mal exposée augmente l'intensité du bruit; des cellules déficientes dans un capteur électronique, inévitables y compris dans le matériel haut de gamme, introduisent un bruit systématique, malgré la tenue d'une liste des pixels défectueux; et ce n'est pas exhaustif.

Dans cette thèse, nous nous intéressons aux images en niveaux de gris, et numériques. La figure 1 montre de telles images contenant des bruits de plusieurs sortes.



FIG. 1 – Bruit naturel dans des images numériques. (a) bruit de sensibilité pour un appareil photo numérique. (b) bruit de scan. (c) bruit de sensibilité + scan. (d) échographie.

Le bruit dans les images est gênant pour plusieurs raisons. Tout d'abord, l'aspect visuel est altéré, la qualité esthétique diminue dans la plupart des cas, et l'œil humain fatigue vite face à des images bruitées. Les logiciels publics de traitement d'images contiennent de plus en plus souvent des fonctions de débruitage; une perspective industrielle est donc en jeu. De plus l'information apportée par une image bruitée est réduite; dans un contexte d'imagerie satellitaire, aérienne ou astronomique où le matériel d'acquisition d'image est cher, chaque donnée est précieuse, et perdre de l'information revient donc à perdre de l'argent. Un autre enjeu est l'amélioration des

algorithmes de détection. En effet, une réduction contrôlée du bruit est un préliminaire important à un algorithme de détection, pour éviter les faux positifs en imagerie médicale par exemple.

Débruiter les images est donc un enjeu important. L'œil humain sait très bien débruiter les images représentant des scènes naturelles, au sens où un bruit d'intensité raisonnable n'empêche pas le système visuel de comprendre une scène, jusqu'à un degré de détail assez fin. Des médecins entraînés arrivent aussi à faire abstraction du bruit quand ils analysent des images médicales. Mais l'automatisation d'une telle tâche n'est pas aisée, parce qu'il s'agit de reconstituer une information perdue pendant le processus de bruitage, sur des données extrêmement variées.

Nous nous intéresserons à un certain type de bruit : nous le supposerons additif (c'est-à-dire qu'il se rajoute aux niveaux de gris), et nous le modéliserons comme la réalisation d'un champ aléatoire blanc et i.i.d. (indépendant et identiquement distribué, c'est-à-dire que le bruit en un pixel est indépendant et a la même loi de probabilité que le bruit en n'importe quel autre pixel), gaussien (la loi du bruit en chaque pixel est gaussienne), et indépendant de l'image originale. Quitte à faire un changement de contraste affine, nous supposerons aussi que la moyenne du bruit est nulle. Autrement dit, si u est l'image originale et ε est un bruit, l'image bruitée v s'écrit

$$v = u + \varepsilon \quad \text{avec} \quad \varepsilon \sim \mathcal{N}(0, \sigma^2 Id).$$

σ^2 représente la variance du bruit, donc en quelque sorte sa puissance. Ce bruit modélise assez bien le bruit de sensibilité ou de scan, même si ce dernier n'est pas exactement i.i.d. La figure 2 montre une image bruitée artificiellement par un bruit blanc gaussien.




FIG. 2 – Ajout artificiel de bruit blanc gaussien (un changement de contraste affine a été appliqué dans l'image de bruit pour le rendre plus visible).

La première méthode historique de débruitage consistait à flouer un peu l'image bruitée, de manière à réduire l'intensité du bruit. Cela peut être exécuté en appliquant une convolution avec un noyau régularisant, ou en faisant évoluer l'image bruitée selon l'équation de la chaleur. Cette méthode permet en effet d'atténuer le bruit, mais a tendance à détruire du même coup de nombreux détails présents dans l'image bruitée. En particulier, les frontières d'objets, qui sont souvent des lieux de discontinuité dans les images naturelles originales, sont lissées, et l'image perd en précision (figure 3 gauche).

Pour éviter de rendre les images floues, [Perona and Malik 1990] ont proposé de lisser l'image bruitée, mais seulement le long des lignes de niveau, de manière à préserver les contours contrastés. Cela s'est traduit par l'introduction d'une équation aux dérivées partielles de diffusion




FIG. 3 – L’image bruitée montrée sur la figure 2 est régularisée par l’équation de la chaleur (gauche) et par une diffusion anisotrope (droite). L’évolution a été arrêtée quand la différence de l’image qui évolue avec l’image initiale bruitée avait la même norme que le bruit, dont la variance était supposée connue.

anisotrope, qui réduit le bruit tout en respectant les contours (figure 3 droite).

Ces techniques reposent sur l’idée que l’image originale est plus régulière que l’image bruitée. Il s’agit donc de rendre l’image la plus régulière possible, mais sans en arriver à effacer des détails importants. Une mesure de régularité d’image peut être un bon outil pour débruiter : il suffirait de minimiser cette quantité, tout en restant suffisamment proche de la donnée.

En 1992, [Rudin et al. 1992] ont proposé une telle mesure de régularité sur les images, la variation totale, qui s’est révélée être un outil particulièrement adapté à de nombreuses thématiques en traitement d’image (réduction du flou, interpolation, désocclusion, réduction des artefacts JPEG, décomposition d’images en structure + texture). La variation totale d’une image numérique u est définie par

$$TV(u) = \sum_x |\nabla u(x)|,$$

où $|\nabla u(x)|$, la norme du gradient de l’image au pixel x , reflète une certaine régularité locale de u au pixel x . La variation totale, somme de ces mesures de régularité locale, estime une certaine forme de régularité globale qui est tout à fait adaptée aux images : elle pénalise fortement les oscillations et les fluctuations aléatoires, tout en autorisant des discontinuités le long de contours suffisamment réguliers. Elle peut de plus s’interpréter comme la longueur cumulée des lignes de niveau de l’image. La minimisation d’une fonctionnelle combinant variation totale et terme d’attache aux données, nommément

$$\lambda TV(u) + \sum_x |v(x) - u(x)|^2 \quad (\text{pour } u \text{ définie sur le même domaine que } v), \quad (1)$$

donne lieu à un débruitage par « minimisation de la variation totale » (minimisation TV), méthode à laquelle nous nous raccrocherons tout au long de cette thèse (cf. figure 4 gauche).

Depuis, la littérature a abondé, à la fois dans l’expansion de méthodes pratiques ou théoriques basées sur la variation totale, et aussi dans la découverte de méthodes de débruitage

bien différentes. Citons notamment les méthodes de seuillage dans la transformée en ondelettes [Mallat 1998], les filtres de voisinage [Yaroslavsky and Eden 1996], et plus récemment, un nouveau filtre non-local [Buades et al. 2006a], appelé « moyennes non-locales » (NL-means). Ces moyennes non-locales sont basées sur des comparaisons de petits bouts d’images (patches), et reconstituent l’information présente dans l’image originale directement à partir de l’information présente dans le contexte bruité (cf. figure 4).




FIG. 4 – L’image bruitée montrée sur la figure 2 est régularisée par une minimisation de la variation totale (gauche) et par les moyennes non-locales (droite). Les paramètres de ces méthodes ont été réglés afin d’avoir, comme en figure 3, une distance quadratique entre l’image bruitée et l’image débruitée égale à la norme du bruit.

Ces méthodes de débruitage ont toutes leurs spécificités. En particulier la minimisation TV et les NL-means sont basées sur des concepts très différents et ont des propriétés bien complémentaires. L’objet de cette thèse est de les mettre en correspondance, malgré leurs différences, grâce à un cadre commun d’estimation statistique bayésienne. C’est dans ce cadre que nous nous plaçons pour former des modèles hybrides : ils sont issus de la minimisation TV, et l’une de leurs caractéristiques est remplacée par ce qu’elle vaut dans le modèle des NL-means.

Trois points principaux séparent la minimisation TV et les NL-means. Tout d’abord le risque bayésien est différent. Dans la première partie de cette thèse nous analysons le débruitage obtenu en adaptant la minimisation TV au risque quadratique des NL-means. Une autre grande différence entre les deux méthodes vient du fait que le critère de régularité dans la minimisation TV est global, tandis que les NL-means n’utilisent que l’information contenue dans un voisinage d’un pixel pour débruiter ce pixel. La localisation du modèle TV est l’objet du premier chapitre de la partie II. Le dernier aspect est le choix du modèle a priori. Dans la minimisation TV, ce modèle est uniquement basé sur la régularité TV, alors que les NL-means exploitent des statistiques exhaustives sur des patches de l’image bruitée. Le choix d’un bon modèle a priori est une question difficile, et nous n’en abordons qu’un petit aspect dans le dernier chapitre.

Le reste de cette introduction est destiné à présenter les chapitres qui suivront.

Partie I : Débruitage TV et moyenne a posteriori

Dans une approche bayésienne, les NL-means sont associées à un risque quadratique, et la minimisation TV à un maximum a posteriori (MAP). Dans la première partie de cette thèse, nous considérons l'adaptation du débruitage par minimisation TV à un risque quadratique. L'image débruitée est alors estimée par la moyenne de la loi a posteriori.

Chapitre 1 : Motivation et définition du débruitage TV-LSE

Le premier chapitre explique les motivations qui sous-tendent ce choix de risque quadratique dans un débruitage par variation totale.

Tout d'abord, la variation totale étant une quantité convexe, la loi a posteriori déduite du modèle est une distribution log-concave, et admet de ce fait une certaine régularité globale. Le point qui réalise le maximum de cette loi indique le lieu de son mode, mais ne tient pas compte d'éventuelles asymétries. L'espérance de la loi a posteriori, elle, est une perturbation du MAP vers les régions de haute densité, qui d'une certaine manière caractérisent mieux la loi a posteriori que le lieu de son mode.

Si la minimisation TV, basée donc sur une estimation MAP, réduit notablement le bruit et préserve les contours, elle subit un artefact bien gênant, l'effet de *staircasing* (marches d'escalier), qui donne aux images débruitées cet aspect constant par morceaux. La présence du phénomène de *staircasing* en débruitage, démontrée par [Ring 2000; Nikolova 1997, 2000, 2004; Malgouyres 2007] dans différents contextes, semble être intimement liée à l'estimation MAP. Très grossièrement, les à-plats dans l'image débruitée correspondent à une norme de gradient nulle, et donc à une variation totale minimale, ce qui les rend très probables. En changeant de risque nous pouvons espérer réduire cet effet de *staircasing*.

L'adaptation à un risque quadratique peut être menée de plusieurs manières. En effet, la minimisation de la variation totale (TV-MAP) peut se faire sous une contrainte égalité de la forme $\|u - v\| = \sigma$ ou une contrainte inégalité $\|u - v\| \leq \sigma$, ou encore sans contrainte à condition d'introduire un multiplicateur de Lagrange (1). Sous des hypothèses peu restrictives, ces trois approches sont équivalentes pour l'estimation MAP, mais se déclinent en trois estimateurs différents pour un risque quadratique. Nous discutons quant à l'équivalence de ces trois estimateurs, notamment en fonction de la dimension de l'image.

Dans la suite de la partie I, nous nous intéressons au débruitage par variation totale avec un risque quadratique, noté TV-LSE (comme *Total-Variation denoising with Least Square Error criterion*), décliné sous sa forme lagrangienne, c'est-à-dire sans contrainte.

Chapitre 2 : Algorithme MCMC pour le débruitage TV-LSE

Une image débruitée par TV-LSE s'écrit comme une certaine intégrale, portant sur l'espace de toutes les images définies sur le même domaine que l'image initiale ; cet espace a une dimension potentiellement très élevée, typiquement de l'ordre de 10^5 . Dans le chapitre 2, nous décrivons un algorithme MCMC (*Monte-Carlo Markov Chain*), ajusté grâce à des travaux récents en statistiques, capable d'approximer efficacement l'estimation TV-LSE.

Grossièrement, il consiste en la réalisation d'une chaîne de Markov d'images dont la loi stationnaire est la loi a posteriori de notre problème. Un théorème ergodique nous fournit la

convergence de la moyenne de Cesàro de la chaîne vers l'image débruitée. Les transitions sont choisies le plus simplement possible, comme un cas de marche aléatoire sur chaque niveau de gris avec acceptation de la transition ou bien rejet, et avec balayage aléatoire des pixels.

La convergence ergodique de la chaîne étant plutôt lente (en $O(\sqrt{N/n})$ où n est le nombre de transitions et N le nombre de pixels dans l'image [Kipnis and Varadhan 1986]), un critère de convergence efficace s'impose. Pour cela nous introduisons une deuxième chaîne de Markov, de même loi mais indépendante de la première. L'écart moyen entre les deux chaînes nous permet de contrôler de manière fine l'erreur commise dans notre approximation.

De plus, pour accélérer la convergence, nous introduisons un paramètre de rodage de la chaîne : on ne commence à moyenniser les itérations de la chaîne qu'à partir d'un certain rang, pour que notre estimation de la moyenne a posteriori ne soit pas trop biaisée par la loi initiale de la chaîne. Le réglage de ce paramètre de rodage, pour un nombre total d'itérations fixé, peut se comprendre comme un compromis biais-variance dans l'estimation de l'image débruitée. Nous décrivons un algorithme simple d'optimisation de ce paramètre.

Un dernier paramètre important pour la vitesse de convergence est le paramètre d'échelle de la marche aléatoire qui est à la base des transitions. L'application de travaux récents de statistiques [Breyer and Roberts 2000; Neal and Roberts 2006] nous permet d'optimiser ce paramètre de manière très simple.

Chapitre 3 : Résultats théoriques

Maintenant que le calcul d'images débruitées par TV-LSE est possible, nous pouvons nous intéresser aux propriétés théoriques attendues de ces images. C'est l'objet du chapitre 3.

Tout d'abord nous nous intéressons à des propriétés élémentaires d'invariance de TV-LSE. Le débruitage TV-LSE hérite des propriétés d'invariance et de symétrie de la variation totale, tout comme la minimisation TV (TV-MAP) : invariance par translation, par rotation du domaine de l'image, mais aussi conservation de la moyenne des niveaux de gris, conservation des propriétés de symétrie de l'image initiale, etc.

La régularité du risque quadratique se transporte aussi sur le débruitage TV-LSE : nous prouvons que l'opérateur de débruitage TV-LSE est continûment différentiable sur l'ensemble des images. C'est la première propriété qui différencie le débruitage TV-LSE de TV-MAP, puisque le débruitage TV-MAP n'est pas partout différentiable.

Nous nous intéressons aussi aux propriétés asymptotiques de TV-LSE, quand les paramètres sont poussés dans un sens ou dans l'autre. Ceci nous permet de situer géométriquement une image débruitée par TV-LSE dans une bande d'espace entre l'image de départ et l'image débruitée par TV-MAP.

Puis nous nous concentrons sur le phénomène de *staircasing*. Nous démontrons que le débruitage TV-LSE n'engendre pas de *staircasing*, contrairement au débruitage TV-MAP. Plus précisément nous prouvons que les zones constantes dans l'image débruitée ont une probabilité nulle. Ce résultat est sans doute l'apport majeur de TV-LSE par rapport au débruitage classique TV-MAP.

La fin du chapitre est consacrée à une étude analytique des débruiteurs basés sur un risque quadratique. Nous considérerons qu'ils sont construits à partir d'une loi a priori plus générale que le modèle TV : nous supposerons la loi a priori seulement log-concave. Des résultats sont

donnés quant à la régularité, la stabilité, et la bijectivité de tels débruiteurs. Ces résultats font pendant aux travaux de [Combettes and Wajs 2004] qui discutent ce genre de propriétés pour les débruiteurs MAP (appelés couramment opérateurs proximaux). Nous démontrons ensuite que sous certaines hypothèses, un débruiteur LSE peut s'écrire sous la forme d'un opérateur proximal, à condition de considérer une loi a priori plus régulière. Ce résultat vient en écho au résultat portant sur l'absence de *staircasing*, puisque [Nikolova 2004] démontre qu'un débruiteur MAP construit sur une loi a priori lisse (C^∞) ne peut générer de *staircasing*. Les conditions d'application de ce résultat ne sont pas totalement réunies dans notre cas puisque nous démontrons seulement que la loi a priori est C^2 , mais cette étude est une première brique de généralisation du résultat d'absence de *staircasing*.

Chapitre 4 : Expériences numériques

Le chapitre 4 regroupe les expériences numériques menées sur le débruitage TV-LSE, utilisant l'algorithme décrit au chapitre 2.

Tout d'abord, nous testons la validité du critère d'arrêt introduit au chapitre 2, basé sur la comparaison de deux chaînes MCMC (*Monte-Carlo Markov Chains*), qui se révèle valide et efficace pour une très large gamme de paramètres, et donne lieu à des temps de calcul raisonnables pour une bonne précision (typiquement 8 minutes pour une image 512×512 pour une précision moyenne de 0.5 niveau de gris sur 255).

Après cela, une discussion sur les deux paramètres principaux du modèle est menée. En particulier, nous montrons que si l'on fait varier l'un de ces paramètres, en réglant l'autre de manière à avoir un niveau de débruitage constant, nous obtenons des images se conformant à un certain compromis entre quantité de *staircasing* et quantité de flou. Le débruitage TV-LSE semble donc remplacer le *staircasing* par un léger flou.

Puis des comparaisons sont menées avec d'autres méthodes classiques de débruitage d'image. Les qualités visuelles de TV-LSE concernent à la fois l'absence de *staircasing*, l'aspect naturel des images débruitées, et le fait que presque aucune structure n'est créée dans du bruit.

Nous considérons aussi une adaptation immédiate du débruitage TV-LSE aux signaux unidimensionnels. Quelques caractéristiques géométriques des signaux originaux, comme leur régularité ou la monotonie, sont beaucoup mieux conservées dans le débruitage TV-LSE que dans le débruitage TV-MAP.

D'autres expériences sur des images montrent que TV-LSE évite aussi le rehaussement de pixels isolés, qui est un autre artefact de TV-MAP ayant lieu pour des faibles niveaux de débruitage, et que TV-LSE reste efficace sur des bruits naturels, ce qui montre que la méthode proposée est plutôt robuste à la modélisation du bruit.

Partie II : Localité et modèle a priori

La première partie de la thèse concerne l'adaptation du débruitage TV au risque intervenant dans les NL-means; nous nous intéressons dans la seconde partie à l'intégration d'autres caractéristiques des NL-means dans le débruitage TV. La première particularité des NL-means (moyennes non-locales) à laquelle nous nous attachons, paradoxalement, est son caractère local. En effet, il a été remarqué notamment par leurs auteurs que le débruitage était beaucoup plus

efficace s'il était localisé, c'est-à-dire si le débruitage d'un pixel ne dépendait que de l'information de l'image contenue dans une certaine fenêtre autour de lui. L'efficacité est double : il diminue beaucoup le temps de calcul, mais surtout, et c'est ce qui nous intéresse ici, donne lieu à un débruitage bien plus homogène, ce qui est visuellement crucial. L'adaptation à un modèle local du débruitage TV est l'objet du chapitre 5. La seconde caractéristique des NL-means est son modèle a priori très riche, basé sur les patches de l'image bruitée. Le chapitre 6 amorce une discussion sur le choix du modèle a priori.

Chapitre 5 : Etude de la localité du débruitage TV

Dans le chapitre 5, nous nous intéressons à la transformation du débruitage TV en un débruiteur local. À la manière des NL-means, le débruitage d'un pixel se fait en débruitant une fenêtre autour de ce pixel selon un certain critère, puis en gardant comme valeur le niveau de gris au centre de la fenêtre. Le débruitage TV devient ainsi un filtre à voisinage. Pour des raisons de temps de calcul, nous nous concentrons sur une version MAP, donnant lieu à un débruitage appelé TV-MAP local, bien qu'une version quadratique semble bien intéressante aussi.

Nous considérons d'abord le cas où chaque fenêtre est débruitée en lui appliquant la version lagrangienne de la minimisation TV. L'adaptation d'un algorithme rapide de minimisation TV [Chambolle 2004] permet des premières simulations, qui se révèlent très mauvaises sur des signaux unidimensionnels. En effet, sur ces signaux, un artefact important intervient, qui ressemble à un effet d'*aliasing* (repliement de spectre). Cet artefact est presque invisible sur les images. Néanmoins, nous proposons une version de TV-MAP local pour les signaux et les images, qui utilise des fenêtres lissées : le terme d'attache aux données est pondéré par la distance au centre de la fenêtre. Certes, ce choix trouve peu de justification en terme de modélisation bayésienne, mais il permet de résoudre totalement ce problème d'« aliasing », et de plus l'idée des fenêtres lissées est aussi utilisée dans les moyennes non-locales, ce qui rapproche un peu plus le débruitage des NL-means. Un algorithme, calqué sur celui de Chambolle, est proposé.

En tant que filtre à voisinage, TV-MAP local devait certainement être associé à une équation aux dérivées partielles. Nous montrons qu'effectivement, sous certaines hypothèses et lorsque la fenêtre est petite, TV-MAP local est équivalent à l'application de l'équation de la chaleur sur un temps très court.

La taille des fenêtres dans le filtre à voisinage est discutée comme un compromis biais-variance. Le lien entre l'optimum de ce compromis et l'échelle caractéristique locale de l'image est discuté, et une tentative d'application, peu concluante, en débruitage adaptatif est présentée.

C'est pourquoi nous proposons une autre version de TV-MAP local capable d'atteindre automatiquement une certaine adaptabilité aux conditions locales. Elle s'appuie sur la minimisation de la variation totale sous contrainte. Ainsi chaque fenêtre débruitée est astreinte à être au plus à une distance donnée de la fenêtre de l'image bruitée correspondante. Cette approche ressemble à celle de [Almansa et al. 2008], où la variation totale est minimisée sous une multitude de contraintes sur le *method noise* local, mais notre méthode est beaucoup plus simple à réaliser, même si elle n'offre pas autant de garanties théoriques.

Chapitre 6 : Variation totale et statistiques locales

Le dernier chapitre est une ébauche de discussion sur le modèle a priori utilisé dans les débruitages par variation totale : le modèle TV.

Dans un premier temps nous construisons une loi a priori sur les images, légèrement différente de la loi de la variation totale, plus cohérente avec les statistiques locales des images naturelles que le modèle TV. Cette loi est déduite des statistiques de la norme du gradient des images naturelles, et est construite grâce à l'algorithme FRAME de [Zhu et al. 1998]. Elle n'est pas sans lien avec des versions régularisées de la variation totale qui donnent des résultats légèrement meilleurs en débruitage. L'existence de cette loi est un élément de réponse à la question suivante : le débruitage TV marche-t-il à cause des propriétés géométriques de la variation totale (formule de la coaire, ...) ou plutôt à cause de ses propriétés statistiques ? En effet, une légère modification de la variation totale la rend plus efficace pour débruiter, alors même qu'elle perd ses propriétés géométriques.

Une autre section est consacrée à l'interprétation du modèle TV et du modèle des NL-means comme les extrêmes d'une décomposition biais-variance dans l'estimation de la distribution a priori, par rapport à la complexité du modèle : la loi de la variation totale est pauvre et induit de gros biais mais de faibles variances, tandis que la loi a priori qui intervient dans les NL-means est basée sur des statistiques exhaustives sur les patches, et donc donne lieu à de faibles biais mais de grandes variances. Ceci laisse à penser qu'un modèle de complexité intermédiaire pourrait atteindre une erreur minimale dans l'estimation de cette loi.

Part I

Total Variation denoising using posterior expectation

Chapter 1

Motivation and definition of the TV-LSE denoiser

Résumé Dans ce chapitre nous proposons un opérateur basé sur la variation totale, pour enlever le bruit additif, blanc et gaussien des images. Il est la transposition de l'opérateur de minimisation de la variation totale, pouvant se voir comme un Maximum A Posteriori (MAP), à un risque quadratique. Nous l'introduisons pour deux raisons. D'abord, les estimées MAP ont tendance à contenir de sérieuses déformations en grande dimension (la dimension est le nombre de pixels dans les images considérées et peut donc être très grande). De plus, la combinaison de la variation totale avec un MAP donne lieu à un effet de *staircasing* (marches d'escalier), et ce dans plusieurs cadres. La méthode proposée s'appuie sur la variation totale (TV) et sur un risque aux moindres carrés (LSE pour *least square error*), nous l'appellerons donc TV-LSE.

Abstract In this chapter, an operator based on Total Variation is proposed to remove additive white Gaussian noise from images. It is a transcription of the Total Variation minimization [Rudin et al. 1992], interpreted as a Maximum A Posteriori (MAP) Bayesian estimation, into a Least Square Error (LSE) risk. It is motivated by two facts. First, MAP estimates are known to behave poorly in high dimension (the dimension, that is the number of pixels in the considered images, can be large). Second, Total Variation combined with MAP gives rise to a staircasing phenomenon in several frameworks. The proposed denoising operator based on both Total Variation and LSE risk, is called TV-LSE.

1.1 Total Variation minimization

When an image has been corrupted by noise, its recovery is not easy, since many images could have lead to the same noisy image. One then needs some kind of regularity criterion to restrict the set of plausible solutions. The Total Variation (TV) is such a regularity measurement, which has proven very efficient in image reconstruction problems. It is well adapted to piecewise constant images (cartoon images), because it allows discontinuities along edges, provided that the edges are located on regular enough curves. Moreover, Total Variation has a nice geometrical

explanation in terms of cumulated length of level lines [Federer 1969; Evans and Garipey 1992; Ambrosio et al. 2000].

Total variation was first applied to image enhancement by Rudin, Osher and Fatemi [Rudin et al. 1992] (see also [Besag 1989] where the *median pixel prior* is proposed, but with no reference to total variation), and then many authors came to use it as an efficient image regularity criterion in variety of topics, such as image recovery [Blomgren et al. 1997; Chambolle and Lions 1997; Chan and Wong 1998; Malgouyres 2002; Combettes and Pesquet 2004], image interpolation [Guichard and Malgouyres 1998], image decompositions [Aujol et al. 2005; Aujol and Chambolle 2005; Aujol et al. 2006; Meyer 2001; Osher et al. 2003], characteristic scale estimation [Luo et al. 2007], JPEG artefacts removing [Alter et al. 2005b]; and this is certainly not exhaustive. This variety of applications shows the power of Total Variation in its ability to model images.

Let us briefly describe how Total Variation broke up in the world of image denoising. Let us begin with an image v , assumed to be a noisy version of an original image u , corrupted by an additive white Gaussian noise ε , i.e.

$$v = u + \varepsilon$$

where u , ε and v are gray level images defined on some measurable subset Ω of \mathbb{R}^2 . We assume v in $L^2(\Omega)$. Our aim is to get u back from v , only assuming that the variance σ^2 of the noise ε is known. Then if the images have large enough dimension, we can expect the empirical variance of the noise to be close to σ^2 (thanks to the weak law of large numbers), and the original image u to lie at a L^2 distance σ from the datum v . We consider then all the images lying in a sphere centered on v with radius σ^2 , and we wonder which one is the best approximation of the original image. We will choose the image which minimizes a certain criterion, a norm for example. The L^2 norm is clearly not adapted to images, because it does not favor any structure or spatial dependence between neighbor pixels, and contrary to human vision, is very sensitive to shifts on the gray levels. We rather take a L^2 norm over the gradient norm of images, which fits human vision more accurately. For this Ω needs to be either an open subset of \mathbb{R}^2 (and we consider images in the Sobolev $W^{1,1}$), either a discrete grid we will assume square, for which the gradient norm is defined by any consistent scheme [Moisan 2007]. Then our estimate for u will be the image which minimizes

$$\int_{\Omega} |\nabla u(x)|^2 dx \quad \text{subject to} \quad \frac{1}{|\Omega|} \int_{\Omega} (v(x) - u(x))^2 dx = \sigma^2,$$

(where $|\Omega| = \int_{\Omega} dx$). The results are numerically disappointing, because the obtained image is blurry. Actually this denoising method is equivalent to a certain thresholding in the Fourier domain, which flattens high frequencies and seriously destroys edges and details.

Rudin, Osher and Fatemi in their seminal paper [Rudin et al. 1992] propose to replace the L^2 norm of the gradient norm by a L^1 norm of the gradient norm, thus retrieving the total variation, defined on the space BV of functions of bounded variation

$$BV(\Omega) = \left\{ u \in L^1(\Omega) \mid \sup_{\varphi \in \mathcal{C}_c^1, |\varphi| \leq 1} \left(\int u \operatorname{div} \varphi \right) < +\infty \right\},$$

by

$$TV(u) = \sup \left\{ \int u \operatorname{div} \varphi, \varphi \in \mathcal{C}_c^1, |\varphi| \leq 1 \right\}.$$

When u lies in $BV(\Omega) \cap C^1(\Omega)$, then $TV(u)$ reduces to

$$TV(u) = \int_{\Omega} |\nabla u|(x) dx.$$

In practice the domain Ω will always have finite mass, and we normalize TV as well as the L^2 norm with respect to $|\Omega|$ the following way

$$TV(u) = \frac{1}{|\Omega|} \int_{\Omega} |\nabla u|(x) dx \quad \text{and} \quad \|u\|^2 = \frac{1}{|\Omega|} \int_{\Omega} (u(x))^2 dx.$$

The denoised image is then chosen to solve the optimization problem

$$\text{minimize } TV(u) \quad \text{subject to } \|v - u\|^2 = \sigma^2. \quad (1.1)$$

This is a constrained formulation of TV minimization with equality constraint. Note that TV is a convex functional but the constraint is not convex. Let us consider the corresponding formulation with inequality constraint

$$\text{minimize } TV(u) \quad \text{subject to } \|v - u\|^2 \leq \sigma^2. \quad (1.2)$$

Here both the functional and the constraint are convex. Let us assume that $\|v - \int v(x) dx\|^2 > \sigma^2$. This is very often true if u and ε are independent random images, and if the image domain is large enough, because then

$$\mathbb{E}\|v - \int v(x) dx\|^2 = \mathbb{E}\|u - \int u(x) dx\|^2 + \mathbb{E}\|\varepsilon - \int \varepsilon(x) dx\|^2 > \text{var } \varepsilon = \sigma^2.$$

Under this assumption, the optimum of (1.2) can be shown to necessarily lie on the boundary of the constraint [Chambolle and Lions 1997], and under this assumption Problems (1.1) and (1.2) are equivalent. A Lagrange formulation for TV minimization will be the most widely used throughout the thesis. It amounts to

$$\text{minimize } \|v - u\|^2 + \lambda TV(u), \quad (1.3)$$

where the parameter λ balances the two terms of the energy, and is an unknown function of σ^2 . Karush-Kuhn-Tucker conditions prove that (1.2) and (1.3) are equivalent for a particular choice of λ .

Let us recall the coarea formula, which is a very nice characterization of TV in terms of cumulated length of level lines. If u is in $BV(\Omega)$, then [Federer 1969; Evans and Gariépy 1992; Ambrosio et al. 2000])

$$TV(u) = \frac{1}{|\Omega|} \int_{-\infty}^{+\infty} \mathcal{H}^1(\partial u_{\lambda}) d\lambda,$$

where u_{λ} denotes the lower level set of u with level λ , ∂u_{λ} denotes its boundary, and \mathcal{H}^1 denotes the 1-dimensional Hausdorff measure. This means that if u is in BV , then its level lines should not be too twisted.

TV minimization has become very popular especially since fast algorithms were found [Hochbaum 2001; Chambolle 2004; Darbon and Sigelle 2005; Goldfarb and Yin 2005; Fu et al. 2006]. It is very efficient when denoising cartoon-like images, i.e. images mainly containing constant regions separated by regular edges. Textures may be erased, because they are regarded as noise. However, flat zones are preserved, and regular edges are well reconstructed. Contrasted and isolated pixels tend to be maintained by the method, often leading to a somehow unnatural aspect. Worse, smooth regions corrupted by noise often suffer from the so-called *staircasing effect*: some gradient intensities are preserved and some others are set down to zero, so that the resulting image is piecewise constant, with artificial boundaries inside a zone which should be smooth.

Methods avoiding texture destruction and staircasing have been the object of recent progress (see [Chan et al. 2000, 2007; Nikolova 2005; Osher et al. 2003] for instance). Essentially their philosophy is to split the image into a texture part and a structure part, to process the texture separately thanks to an appropriate method, and to denoise the structure part with methods derived from TV minimization sometimes with a smoothed version of the total variation, which then loses its geometrical interpretation in terms of level lines.

However it seems that the staircasing artefact is not caused by the image model itself defined by Total Variation, but precisely by the variational framework in which it lies.

1.2 TV minimization and Maximum a Posteriori

1.2.1 A statistical framework for TV minimization

It is a well-known fact that TV minimization can be understood as a Maximum a Posteriori (MAP) estimation. Let us explain why here.

First let us assume that the considered images are defined on a finite domain $\Omega \subset \mathbb{R}^2$. For a given noisy datum v , let E_λ denote the Rudin-Osher-Fatemi (ROF) energy [Rudin et al. 1992]

$$E_\lambda(u) = \|u - v\|^2 + \lambda TV(u)$$

which must be minimized when solving Problem (1.3). Let now \hat{u} denote the solution of this variational formulation, that is

$$\hat{u} = \arg \min_u E_\lambda.$$

Let us explain why \hat{u} can be viewed as a Maximum a Posteriori (MAP) estimate of the original image u . Let $\beta > 0$ and p_β the prior probability density function (*TV prior*) defined by

$$p_\beta(u) = \frac{1}{Z} e^{-\beta TV(u)}$$

where Z is a universal normalizing constant (see the remark below concerning the well-posedness of p_β). This prior distribution models the gradient norms on each pixel as independent and identically distributed random variables following a Laplace distribution. This model does not exactly fit the reality, because very often the high gradient norms are concentrated along curves for instance, and hence are not independent. Furthermore it assumes that realistic images have finite total variation, which was found to be untrue [Gousseau and Morel 2001]: in a multiscale approach, the total variation is expected to blow up to infinity when the image resolution

increases. However, this model is convenient and efficient for our tasks, and in the sequel we shall see that it is even better than could be expected.

We assumed the noise to be Gaussian, white and additive. Then the conditional density of v given u , i.e. the degradation density, is given by

$$p(v|u) = \frac{1}{Z} e^{-\frac{\|v-u\|^2}{2\sigma^2}}.$$

Now that a prior density and degradation density have been defined, Bayes' rule can be applied to the denoising of image v , and then the posterior density of u given v is

$$p_\beta(u|v) = \frac{p(v|u)p_\beta(u)}{\int p(v|u')p_\beta(u')du'} = \frac{1}{Z} \exp \left[- \left(\frac{\|v-u\|^2}{2\sigma^2} + \beta TV(u) \right) \right],$$

where Z only depends on v . The maximization of the posterior density is equivalent to the minimization of the Rudin-Osher-Fatemi energy, i.e.

$$\text{maximize } p_\beta(u|v) \iff \text{minimize } E_\lambda,$$

provided that $\lambda = 2\sigma^2\beta$. Then, image denoising by TV minimization is merely a Bayes estimation, associated to a Maximum A Posteriori (MAP).

Remark 1: Actually the prior p_β is not well defined, because for any fixed positive value r ,

$$TV(u) = r \implies \forall c \in \mathbb{R}, TV(u+c) = r,$$

and this implies that $e^{-\beta TV(u)}$ is not integrable on the whole space of images \mathbb{R}^Ω . The prior p_β is referred to as an *improper* prior [Besag 1989]. To overcome this problem, we can reduce the space of integration. We briefly present two ways to achieve this. First of all let us introduce for all $m \in \mathbb{R}$ the subspace

$$\mathcal{E}_m = \left\{ u \in \mathbb{R}^\Omega \mid \frac{1}{|\Omega|} \sum_{x \in \Omega} u(x) = m \right\},$$

which is the hyperplane of images whose gray level average is m . If p_β is restricted on the images of $\mathcal{E}_{\bar{v}}$ (where \bar{v} is the average of v), then p_β is properly defined as soon as it is integrable on $\mathcal{E}_{\bar{v}}$ (the integrability may depend on the numerical scheme for TV). It is consistent to consider such a restriction in high dimension since with high probability the original image u has an average \bar{v} thanks to the weak law of large numbers. This defines a first prior (this restriction was applied in the original paper [Rudin et al. 1992], although there was no reference to MAP estimation).

Now we can come closer to the primitive definition of p_β by introducing a prior on the image's average. Let $\pi(m)$ be a probability density function modeling such a prior on the images' average. Then a second consistent total variation prior density on images could be

$$p_\beta(u) = \frac{1}{Z} e^{-\beta TV(u)} \pi(\bar{u}).$$

For example, letting $\pi_M(m) = \frac{1}{2M} \mathbf{1}_{-M < m < M}$ (which amounts to consider a uniform measure on $\cup_{-M < m < M} \mathcal{E}_m$), one has

$$p_\beta(u) = \frac{1}{Z} e^{-\beta TV(u)} \mathbf{1}_{-M < \bar{u} < M}.$$

Taking M large enough consistently approximates the total variation distribution, and the posterior distribution with density

$$\frac{1}{Z} \exp\left(-\frac{\|v - u\|_2^2}{2\sigma^2}\right) e^{-\beta TV(u)} \pi_M(\bar{u})$$

converges towards the well-defined distribution with density

$$\frac{1}{Z} \exp\left(-\frac{\|v - u\|_2^2}{2\sigma^2}\right) e^{-\beta TV(u)}$$

when M goes to $+\infty$.

Remark 2: A probability density function p is said to be logarithmically concave (log-concave) if its potential $-\log p$ is convex (a more general definition of log-concavity can be found in [Prékopa 1973]). Here p_β is log-concave because its potential βTV is convex. The posterior distribution $p_\beta(u|v)$ is also log-concave. Log-concavity implies a certain global regularity for the probability density function (p.d.f.). In particular, it implies unimodality (a unique maximum).

In the sequel we shall omit the index β in the prior $p_\beta(u)$ and in the posterior $p_\beta(u|v)$.

1.2.2 Loss functions

From a general posterior distribution $p(u|v)$, not necessarily derived from Total Variation, the MAP estimate is the maximizer of $p(u|v)$. This choice is natural, but of course most information on $p(u|v)$ is being lost. For example the reliability of this estimate is lost, although $p(u|v)$ gives some information about it, via the existence of other local maxima or the posterior's curvature at the global maximum for instance. The integration of posterior information into the estimate can be considered via the introduction of loss functions.

Let L be a loss function, i.e. a mapping $L : \mathbb{R}^\Omega \times \mathbb{R}^\Omega \rightarrow \mathbb{R}$, where $L(u, u')$ measures the cost of estimating u' by u instead of the true data u' . Three typical examples of loss functions are the following (see Figure 1.1)

$$\begin{cases} L_0(u, u') &= \begin{cases} 0 & \text{if } u = u' \\ 1 & \text{otherwise} \end{cases} \\ L_1(u, u') &= \|u - u'\|_2 \\ L_2(u, u') &= \|u - u'\|_2^2 \end{cases}$$

The L_0 loss is a “hit-or-miss” loss which equally penalizes any wrong estimate, whereas L_1 and L_2 are more regular costs which penalize distant estimates more than closer ones. Other loss functions have been considered in the context of image processing (see for instance [Rue and Hurn 1997]), but here we only focus on the simplest ones.




Figure 1.1: Classical loss functions on \mathbb{R} . (a) “Hit-or-miss” loss function: f_0 is plotted, where $L_0(u, u') = \delta_{u \neq u'} = f_0(u - u')$. (b) Homogeneous loss function: f_1 is plotted, where $L_1(u, u') = f_1(u - u') = \|u - u'\|_2$. (c) Quadratic loss function: $L_2(u, u') = f_2(u - u') = \|u - u'\|_2^2$.

Let $p(u)$ be a prior distribution on \mathbb{R}^Ω , and $p(v|u)$ be the degradation distribution; then the posterior distribution $p(u|v)$ is known. The risk R associated to a loss L is defined as the expectation of the loss function with respect to u' following the posterior distribution

$$R(u) = \mathbb{E}_{u'|v}[L(u, u')] \doteq \int_{\mathbb{R}^n} L(u, u') p(u'|v) du'.$$

A general framework in Bayes estimation is to consider as an estimate \hat{u} the image which minimizes this risk, which is

$$\hat{u} = \arg \min_u \mathbb{E}_{u'|v}[L(u, u')].$$

Let us derive this model with the three particular losses mentioned above.

The L_0 loss gives rise to the risk

$$R_0(u) = \mathbb{E}_{u'|v}[1 - \delta_{u=u'}] = \frac{1}{Z} \int \mathbf{1}_{u \neq u'} dp_{u'|v} = 1 - \frac{1}{Z} \int \mathbf{1}_{u=u'} dp_{u'|v}.$$

Minimizing this risk is selecting the image u which will be equal to u' when $p_{u'|v}$ is maximal. This exactly corresponds to MAP estimation. Indeed,

$$\arg \min R_0(u) = \arg \max_{u'} p_{u'|v} = \hat{u}_{MAP}.$$

The estimate \hat{u}_{MAP} is the most probable image given v , i.e. the image which maximizes the posterior distribution. Hence TV minimization corresponds to Bayes estimation with a “hit-or-miss” loss function.

Now the risk derived from the L_1 loss is

$$R_1(u) = \mathbb{E}_{u'|v} \|u - u'\|.$$

The minimum of this function of u is reached at the median of the posterior distribution

$$\arg \min \mathbb{E}_{u'|v}[L_1(u, u')] = \text{med} p_{u'|v} = \hat{u}_{med}.$$

Last but not least, the minimal risk derived from the L_2 loss is reached at the expectation of the posterior distribution. We will speak of it as the Least-Square Error (LSE) estimate of u

$$\arg \min \mathbb{E}_{u'|v}[L_2(u, u')] = \mathbb{E}[u|v] = \hat{u}_{LSE}.$$

To our knowledge, posterior expectation has never been applied jointly with a TV prior distribution to image denoising. The next subsection is devoted to several problems encountered in MAP estimation combined with image denoising. This will motivate us to transpose the TV minimization into a LSE framework.

1.2.3 Distortions of MAP estimation in image denoising

Distortion due to high dimension

The MAP estimate is a biased summary of the posterior. In particular, typical images drawn from a high-dimensional distribution are far from maximizing the density. Moreover, according to [Rue and Hurn 1997], the MAP estimate only depends on the location of the mode, but not on the probability mass this mode contains; and this difference becomes potentially huge when working in high dimension. This is still true for log-concave densities considering the mass contained in a neighborhood of the mode's maximizer. This is illustrated by the following proposition.

Proposition 1.1 *Let $n \in \mathbb{N}^*$. Let X be a random vector distributed as $\mathcal{N}(0, \sigma^2 I_n)$ (normal centered distribution with covariance matrix $\sigma^2 I_n$). Let $\|X\|$ be its dimension-standardized norm, i.e. such that $\|X\|^2 = \frac{1}{n} \sum_{i=1}^n X_i^2$. Then*

$$\forall \varepsilon > 0, \quad \mathbb{P}(|\|X\|^2 - \sigma^2| > \varepsilon) \leq \frac{2\sigma^4}{n\varepsilon^2}$$

which tends towards 0 when the dimension n goes to infinity.

Proof : This is merely a consequence of Bienaymé-Tchebychev inequality

$$\forall \varepsilon > 0, \quad \mathbb{P}(|\|X\|^2 - \mathbb{E}[\|X\|^2]| > \varepsilon) \leq \frac{\text{var}\|X\|^2}{\varepsilon^2}.$$

As $n\|X\|^2$ follows a $\text{sigma}^2 \chi^2(n)$ distribution, $\|X\|^2$'s mean is σ^2 and its variance is $\frac{2\sigma^4}{n}$. This yields the result. \square

In this proposition, the mode of X is located at 0, but when n goes to $+\infty$ all the mass of the distribution “concentrates” onto the sphere centered at 0 with radius $n\sigma$, and therefore goes away from the mode. Estimation where the optimization is based on pointwise densities (which is done by MAP estimation) therefore seems questionable in high dimension.

We give here another simple proposition illustrating the fact that in high dimension the major part of the probability mass concentrates on the border of the probability support.

Proposition 1.2 *Let $n \in \mathbb{N}^*$. Let $X \sim U_{B_n(0,1)}$ be a random image uniformly distributed on the unit ball of \mathbb{R}^n . Let $\varepsilon > 0$ and $X'_\varepsilon \sim U_{B_n(0,1) \setminus B_n(0,1-\varepsilon)}$ uniformly distributed on the centered crown with radii $1 - \varepsilon$ and 1 . Then the probability distributions p_X and $p_{X'_\varepsilon}$ of X and X'_ε satisfy*

$$d_{TV}(p_X, p_{X'_\varepsilon}) := \sup_{A \text{ Borel set}} |p_X(A) - p_{X'_\varepsilon}(A)| \xrightarrow{n \rightarrow \infty} 0$$

Proof : The supremum is attained for $A = B_n(0, 1 - \varepsilon)$, which yields

$$d_{TV}(p_X, p_{X'_\varepsilon}) = \frac{|B_n(1 - \varepsilon)|}{|B_n(1)|} = (1 - \varepsilon)^n \frac{|B_n(1)|}{|B_n(1)|} = (1 - \varepsilon)^n \xrightarrow{n \rightarrow \infty} 0,$$

where $|A|$ denotes the volume of A . □

Again all the mass of the distribution escapes to the outskirts of the support space, essentially because in high dimension the volume of the unit ball is “concentrated” on its border. This corroborates the fact that the MAP estimate is a very special image, whose features really differ from those of typical samples of the posterior. The following paragraphs demonstrate that MAP estimation can lead to serious artefacts when applied to image denoising.

Distortions in simple cases of MAP estimation

The Bayesian approach consists in modeling parameters (in our case the true image u) as random variables following a prior distribution, before using the posterior distribution in the estimation.

The modeling of the parameter distribution can be thought of as a prior design problem. The prior is often built from experimental considerations and translates u 's natural behavior as faithfully as possible into a probability language.

In particular, one could expect that the estimate \hat{u} , for instance, should be a typical simulation of the prior distribution p_u . One could additionally expect that the noise estimate $v - \hat{u}$ should ideally look like a simulation of a white Gaussian noise. This last idea is common in image denoising: people find it nice when the noises induced by the method is white (i.e. contains no structure in practice). Actually the two requirements are dual thanks to the symmetry between the prior distribution and the noise distribution, described in Proposition 1.3.

Proposition 1.3 *Let p_U and p_ε be distributions on U and ε respectively. Let L be a radially symmetric risk, h i.e. such that*

$$\exists f : \mathbb{R} \rightarrow \mathbb{R}, \quad L(u, u') = f(\|u - u'\|).$$

Then the two estimates of u

$$\begin{cases} \hat{u} = \arg \min_u \mathbb{E}_{u'|v}[L(u, u')] \\ \hat{u}' = v - \arg \min_\varepsilon \mathbb{E}_{\varepsilon'|v}[L(\varepsilon, \varepsilon')] \end{cases}$$

are identically equal, as well as the two estimates of ε

$$\begin{cases} \hat{\varepsilon} = v - \arg \min_u \mathbb{E}_{u'|v}[L(u, u')] \\ \hat{\varepsilon}' = \arg \min_\varepsilon \mathbb{E}_{\varepsilon'|v}[L(\varepsilon, \varepsilon')]. \end{cases}$$

Proof : Making the change of variable $u = v - \varepsilon$ and $u' = v - \varepsilon'$ leads to

$$\begin{aligned}\hat{u}' &= v - \arg \min_{\varepsilon} \mathbb{E}_{\varepsilon'|v}[L(\varepsilon, \varepsilon')] \\ &= v - \left[v - \arg \min_u \mathbb{E}_{u'|v}[L(v - u, v - u')] \right] \\ &= \arg \min_u \mathbb{E}_{u'|v}[L(u, u')],\end{aligned}$$

which proves $\hat{u} = \hat{u}'$. Now the equalities $\hat{\varepsilon} = v - \hat{u}$ and $\hat{\varepsilon}' = v - \hat{u}'$ imply $\hat{\varepsilon} = \hat{\varepsilon}'$. \square

Unfortunately such priors seldom occur, if ever. Nikolova [2007] gives theoretical explanations and numerous examples for which the MAP estimate deviates from the prior distribution, in the case of single parameter estimation and discrete signals. Indeed, in the models studied, both the object estimate and noise estimate follow a strongly biased distribution, which give them properties the simulations from the prior seldom have.

For example, let the prior p_u defined on scalar random variables (for images this amounts to consider independently distributed pixels, which more or less occurs after a wavelet transformation e.g. [Buccigrossi and Simoncelli 1999]), by $p_u(u) = \frac{1}{Z}e^{-\lambda|u|^\alpha}$, and let us still consider a Gaussian noise $p_\varepsilon(\varepsilon) = \frac{1}{Z}e^{-\varepsilon^2/(2\sigma^2)}$. As shown by [Nikolova 2007], for $\alpha \geq 1$, the distribution of the estimate \hat{u} is more strongly peaked than the prior in the numerical experiments. The distribution of the noise estimate $v - \hat{u}$ is supported by a bounded interval, and is strongly peaked on the bounds of this interval: this is very far from a Gaussian distribution. For $0 < \alpha < 1$, the distribution of \hat{u} is essentially concentrated on zero, while $v - \hat{u}$ again is bounded.

Another mentionable example of such a distortion is also developed in [Nikolova 2007] and illustrated on Figure 1.2. A discrete signal u is modeled as a random walk on \mathbb{R} with Laplacian transition, i.e. the shifts $(u_{i+1} - u_i)_i$ are independent and identically distributed (i.i.d.) and follow a distribution of the form $\frac{\lambda}{2}e^{-\lambda|u_{i+1} - u_i|}$ (for short, the resulting signal is a Markov chain with Laplace increments). A white Gaussian noise is again assumed. Note that the MAP denoising of u with such assumptions is the minimization of the total variation of u transposed to signals with the particular numerical scheme for total variation

$$TV(u) = \sum_i |u_{i+1} - u_i|.$$

It is proven that the probability of having $(\hat{u}_{i+1} - \hat{u}_i) = 0$ for a fixed i is non zero. Numerical experiments amplify this reality: the distribution of $(\hat{u}_{i+1} - \hat{u}_i)$ is mainly a huge peak at zero, with 92% of mass inside this peak, unlike the Laplacian prior which emphasizes to the non-zero values. It results in a piecewise constant estimate \hat{u} , which is visually very far from a simulation of the signal according to the prior.

[Nikolova 2007] concludes by this striking remark: this model is much better to denoise signals from *impulse noise* than from Gaussian noise. Then a prior design (of noise model, but also of u thanks to the duality) based on natural behavior seems to be a rough approximation only of the best prior, in these cases of MAP estimation. Indeed, Bayesian estimation is relevant in a modeling point of view, but offers little guarantee about the quality of the result. One could think of modifying the prior and the degradation models until the MAP estimate becomes relevant

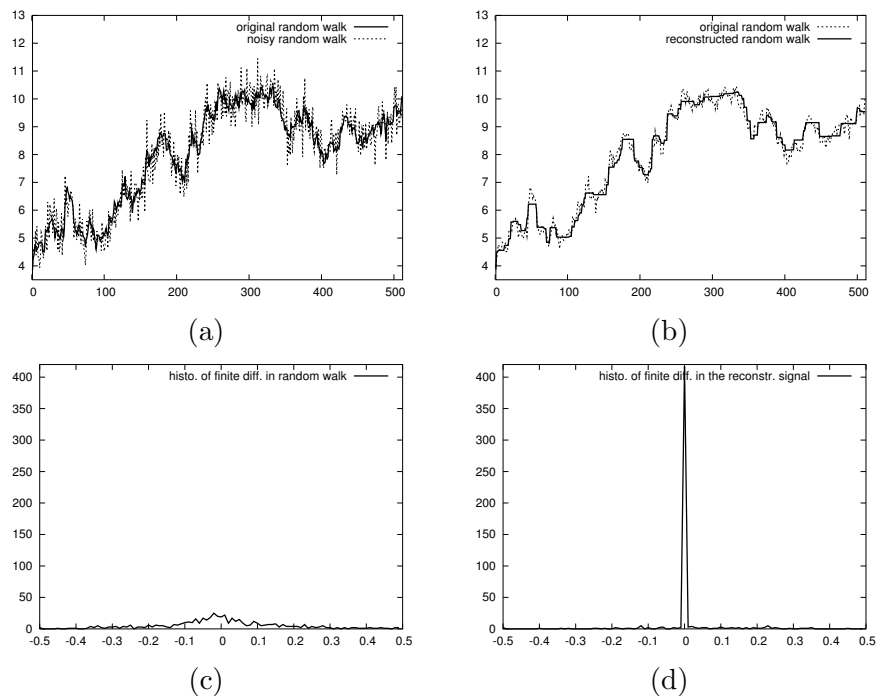


Figure 1.2: Illustration for 1-D staircasing. (a) (u_n) is a random walk (solid curve) with Laplacian transition (i.e. $(u_{n+1} - u_n)$ i.i.d. $\sim \frac{\lambda}{2}e^{-\lambda|x|}$ with $\lambda = 8$) (bold), and its noisy version (dashed curve) $(u_n + \varepsilon_n)$ with ε_n i.i.d. $\sim \mathcal{N}(0, \sigma^2)$ where $\sigma = 0.5$ is assumed to be known. (b) The MAP estimate (\hat{u}_n) (bold) and again the original (u_n) . Many regions are estimated as constant regions when it should not be the case. This illustrates the 1-D staircasing effect. (c) Histogram of the transitions $(u_{n+1} - u_n)$. It follows a Laplacian distribution. (d) Histogram of the estimated transitions $(\hat{u}_{n+1} - \hat{u}_n)$ with same scale as in (c). The distribution is impressively peaked at zero, and this peak contains more than 80% of the total mass. This peak is due to the ubiquitous staircasing effect in the estimated chain (\hat{u}_n) . These graphs were simulated back from [Nikolova 2007].

(somehow this approach is the object of L^1 and L^0 minimization problems [Fu et al. 2005; Tropp 2006], $u + v$ decompositions using G or E spaces [Aujol et al. 2005; Aujol and Chambolle 2005; Aujol et al. 2006]). But another approach would simply be to try loss functions other than the “hit-or-miss”.

Staircasing for TV-denoising

The previous example where the estimates are mostly piecewise constant exhibits the *staircasing* artefact induced by TV-denoising on discrete signals. An archetype of staircasing effect can easily be rendered only by taking a linear signal defined on a discrete segment, by adding noise to it and to denoise it by total variation minimization. Results are shown on Figure 1.3. This artefact can also be found in images denoised by TV minimization, as in Figures 1.4 and 1.5 .

The word *staircasing* intuitively means two things: first the existence of constant regions, but also the existence of discontinuities. Now, discontinuities were already analyzed in detail




Figure 1.3: Staircasing on signals. (a) A non-noised affine ramp (dashed curve) denoised by TV minimization (solid curve): apart from the borders, the denoised signal is faithful to the datum. (b) An affine ramp modulated by a sine curve (dashed), and its denoised version (solid): the singularities occur at the maximal derivative of the datum. There is staircasing but no discontinuity. (c) A noised affine ramp (dashed) and its denoised version (solid): discontinuities occur at many locations together with creation of constant zones. These three graphs have been simulated back from [Caselles et al. 2007].

in [Caselles et al. 2007]. Furthermore, Figure 1.3 (b) shows a signal with no discontinuity, but visually containing serious staircasing artefact. This means that the discontinuities may be less crucial for human vision than the existence of constant zones. This justifies the following definition for staircasing.

We define the connectedness of a set $\Omega_1 \subset \Omega$ by the 4-pixel neighborhoods. We will speak of *staircasing effect* in the denoised version \hat{u} of a noisy image if there exists at least one connected region $\Omega_1 \subset \Omega$ containing at least two pixels, such that

$$\forall x \in \Omega_1, \quad \forall x' \in \mathcal{N}_x, \quad x' \in \Omega_1 \Rightarrow \hat{u}(x) = \hat{u}(x'),$$

where \mathcal{N}_x denotes a neighborhood of pixel x (both in the discrete and continuous settings). Of course, staircasing becomes visually unpleasant when the sets Ω_1 are big and many. Figures 1.4 and 1.5 show this artefact on images denoised by TV minimization. It becomes more and more visible when λ is large.

Staircasing effect was first pointed out by [Dobson and Santosa 1996]. They made use of this artefact to reconstruct “blocky” images. Afterwards, [Blomgren et al. 1997; Chan et al. 2000, 2007] proposed to avoid staircasing effect in natural images by different methods. Namely several authors [Bouman and Sauer 1993; Blomgren et al. 1997] propose to replace the total variation by a functional of the form

$$\sum_x f(|Du|(x)) \text{ where } f(t) \sim \begin{cases} t^2 & \text{near } 0 \\ t & \text{near } +\infty. \end{cases}$$

Let us also cite [Buades et al. 2006b; Savage and Chen 2006] for a short overview of staircasing reduction methods in image processing.

In the meantime authors mathematically proved the existence of this staircasing effect, in the one-dimensional continuous case [Ring 2000] as well as for the two-dimensional discrete case [Nikolova 1997, 2000, 2004, 2005]. Namely in [Nikolova 2004] the total variation is viewed as a




Figure 1.4: Staircasing on images (I). (a) noisy image of Lena (standard deviation of the noise $\sigma = 10$). (b) the image denoised by TV minimization ($\lambda = 30$). (c) and (d): details. Constant regions are created, separated with spurious edges, when the original image is smooth. This artefact is known as the staircasing effect. (e) and (f): the level lines (bilinear interpolation) concentrate on spurious edges.

particular case of operators with the form

$$J : u \mapsto \sum_{x=1}^r \varphi(G_x(u)),$$

where G_x are linear operators, and φ is a non-differentiable at zero function, not necessarily convex. [Nikolova 2004] states the following result under quite loose conditions.

Proposition 1.4 [Nikolova 2004] *If $\mathcal{X}(v) = \arg \min_u \|u-v\|^2 + \lambda J(u)$ where $J(u) = \sum_{i=1}^r \varphi_i(G_i u)$, with G_i linear and φ a non-differentiable at 0 potential function, then there exists a neighborhood V_L of v for which*

$$\forall v' \in V_L, \{i \mid G_i(\mathcal{X}(v')) = 0\} = \{i \mid G_i(\mathcal{X}(v)) = 0\}. \quad (1.4)$$




Figure 1.5: Staircasing on images (II). (a) noisy image of Barbara (standard deviation of the noise $\sigma = 10$). (b) the image denoised by TV minimization ($\lambda = 30$). (c) and (d): details. Again we can see staircasing effect. (e) and (f): corresponding level lines (bilinear interpolation).

In particular, the case where J is a Total Variation operator such as

$$TV(u) = C \sum_{x \in \Omega} \sum_{\|y-x\|=1} |u(y) - u(x)|$$

fulfills the conditions of Proposition 1.4, where the (G_x) stand for first-order differences between pixel x and its neighboring pixels, and φ stands for the absolute value function, which is not differentiable at 0. In this case, the set $\{i \mid G_i(\mathcal{X}(v)) = 0\}$ contains the regions where the denoised image is constant. The proposition tells that the constant regions of denoised images have a certain stability with respect to the noise. This gives a first theoretical explanation of the staircasing effect.

Let us cite the recent work of Caselles, Chambolle and Novaga in [Caselles et al. 2007] where the point of view on staircasing does not concern constant regions here but rather discontinuities. An interesting property concerning the jump set of the reconstructed image in the continuous framework is proven, which could suggest that staircasing is only due to a bad quantization of

the total variation. The (approximate) jump set of a continuous image u is defined as the set of points x satisfying

$$\exists u^+(x) \neq u^-(x), \exists \nu_u(x) \in \mathbb{R}, \begin{cases} |\nu_u(x)| = 1, \\ \lim_{\rho \downarrow 0} \frac{\int_{B_\rho^+(x, \nu_u(x))} |u(y) - u^+(x)| dy}{\int_{B_\rho^+(x, \nu_u(x))} dy} = 0, \\ \lim_{\rho \downarrow 0} \frac{\int_{B_\rho^-(x, \nu_u(x))} |u(y) - u^-(x)| dy}{\int_{B_\rho^-(x, \nu_u(x))} dy} = 0, \end{cases}$$

where

$$B_\rho^+(x, \nu_u(x)) = \{y \mid \|y - x\| < \rho, \langle y - x, \nu_u(x) \rangle > 0\}$$

and $B_\rho^-(x, \nu_u(x))$ is the same with a negative scalar product. Intuitively it is the set of points where the image u can be locally described as a two-dimensional Heaviside function. This corresponds to regular edges. It is shown that if the datum image v has bounded variation, then the jump set of the solution \hat{u} to the TV minimization problem is contained within the jump set of v . In other words, TV minimization does not create edges which did not already exist in v . This would contradict some kind of staircasing effect (the discontinuity part), if we forgot that v is generally noisy and then the jump set contains almost every point of the domain.

Temperature problem

Consider the following temperature problem (it can be transposed to any scalar estimation problem). Imagine that meteorological services provide a probability distribution of tomorrow's temperature, which is not a Dirac, because there is uncertainty about it. The question is: what temperature would you bet for? In other words, what is the best way to compress the information brought by a complete probability distribution into a single real value? Here we discuss this question, but do not pretend to solve it.

If the temperature distribution is symmetric and unimodal (see Figure 1.6 left) then a natural choice would be the maximizer of the distribution, and this exactly corresponds to the expectation (LSE) and to the median estimates.

Now if the temperature distribution has two peaked modes (see Figure 1.6 center), what would be the best estimate? A temperature corresponding to one of the peaks, or the point in the center? Actually, none is satisfying: choosing one peak among the two is an arbitrary choice which is not satisfactory, while the center between the two peaks is very unlikely to occur.

Last but not least, consider the case where the temperature is unimodal but asymmetrical (see Figure 1.6 right). The MAP estimate corresponds to the maximizer of the temperature distribution, while the LSE estimate as well as the median estimate correspond to a shift of the MAP estimate towards the heavier side of the distribution. The median as well as the LSE estimates seem to be good trade-offs in terms of "compression" of the temperature distribution.

The case of TV minimization corresponds to the third case, in the sense that the posterior distribution is log-concave (i.e. $-\log p(u|v) = \frac{1}{2\sigma^2} E_\lambda(u) - \log Z$ is convex) and thus unimodal, and asymmetrical. That is why a median and a LSE estimates are likely to better compress the information given by the posterior, by shifting the MAP into the direction of greater mass of the posterior.




Figure 1.6: 3 particular cases of scalar posterior (temperature) distribution. Left: the distribution is symmetric and unimodal; then the MAP equals the median and the mean. Any other estimate would not correspond to the regularity properties of the posterior. Center: the distribution is bimodal and symmetric; the MAP will fetch arbitrarily one of the two maxima, whereas the median and the mean correspond to the local minimum (the symmetry center). None of these choices is satisfying, because a MAP estimate would not respect the symmetry property of the posterior, and both the mean and the median would correspond to a risky choice. Right: the distribution is unimodal but asymmetrical. The median and the mean correspond to a shift of the MAP towards the heavier tail of the distribution. Shifted versions (as mean and median) have the property to have maximum total mass inside a neighborhood of the estimate, for a certain size of neighborhood ($= 0$ for the MAP, > 0 for mean and median). These shifted estimates can be of interest depending on the application, in image denoising in particular.

1.3 Definition of TV-LSE

In this section at last, we introduce a denoising estimation based on the TV prior and the posterior expectation. Compared to TV minimization, only the loss function has been changed from the “hit-or-miss” loss function L_0 into the quadratic (“least-square error”) loss function L_2 . Such an estimator is given by the expectation of the posterior (also known as *posterior mean* [Besag 1989])

$$\hat{u}_{LSE} = \int_{\mathbb{R}^{\Omega}} u p(u|v) du = \frac{\int_{\mathbb{R}^{\Omega}} u e^{-\frac{\|u-v\|^2 + \lambda TV(u)}{2\sigma^2}} du}{\int_{\mathbb{R}^{\Omega}} e^{-\frac{\|u-v\|^2 + \lambda TV(u)}{2\sigma^2}} du}.$$

Proposition 1.5 \hat{u}_{LSE} is well defined.

Proof : The denominator satisfies

$$0 < \int e^{-\frac{\|u-v\|^2 + \lambda TV(u)}{2\sigma^2}} du \leq \int e^{-\frac{\|u-v\|^2}{2\sigma^2}} du = \prod_{i \in \Omega} \int_{\mathbb{R}} e^{-\frac{(u_i - v_i)^2}{2\sigma^2}} du_i < +\infty,$$

and the numerator at pixel i satisfies

$$\int |u_i| e^{-\frac{\|u-v\|^2 + \lambda TV(u)}{2\sigma^2}} du \leq \int |u_i| e^{-\frac{\|u-v\|^2}{2\sigma^2}} du = \int |u_i| e^{-\frac{(u_i - v_i)^2}{2\sigma^2}} du_i \prod_{j \in \Omega, j \neq i} \int_{\mathbb{R}} e^{-\frac{(u_j - v_j)^2}{2\sigma^2}} du_j < +\infty.$$

The integrals are absolutely convergent, hence \hat{u}_{LSE} is well defined. \square

Remark: \hat{u}_{LSE} depends on two parameters λ and σ (or equivalently on $\beta = \frac{\lambda}{2\sigma^2}$ and σ) while \hat{u}_{MAP} only depends on λ . In other words, \hat{u}_{MAP} merges the desired regularity β and the presumed quantity of noise σ into a single parameter $\lambda = 2\sigma^2\beta$, while \hat{u}_{LSE} allows the distinction of these two parameters.

1.3.1 Constrained TV-LSE

Let us denote $B(v, \sigma)$ the L^2 ball centered at v with radius σ . The constrained formulation of TV minimization

$$\begin{aligned} \hat{u}_{MAP} &= \arg \min_u TV(u) \quad \text{s. t.} \quad \|u - v\| \leq \sigma \\ &= \arg \min_u TV(u) + \chi_{B(v, \sigma)}(u) \quad \text{where } \chi_{B(v, \sigma)}(u) = \begin{cases} 0 & \text{if } \|u - v\| \leq \sigma \\ +\infty & \text{otherwise,} \end{cases} \end{aligned}$$

can also be adapted to the LSE risk, yielding a new formulation of TV-LSE

$$\begin{aligned} \hat{u}_{LSE}^{(B)} &= \frac{1}{Z} \int u e^{-\beta[TV(u) + \chi_{B(v, \sigma)}(u)]} du \\ &= \frac{\int_{B(v, \sigma)} u e^{-\beta TV(u)} du}{\int_{B(v, \sigma)} e^{-\beta TV(u)} du}. \end{aligned}$$

Likewise, letting $S(v, \sigma)$ the sphere centered at v with radius σ , the other constrained formulation of TV minimization

$$\begin{aligned} \hat{u}_{MAP} &= \arg \min_u TV(u) \quad \text{s. t.} \quad \|u - v\| = \sigma \\ &= \arg \min_u TV(u) + \chi_{S(v, \sigma)}(u) \quad \text{where } \chi_{S(v, \sigma)}(u) = \begin{cases} 0 & \text{if } \|u - v\| = \sigma \\ +\infty & \text{otherwise,} \end{cases} \end{aligned}$$

can be transposed to LSE risk, yielding

$$\hat{u}_{LSE}^{(S)} = \frac{\int_{S(v, \sigma)} u e^{-\beta TV(u)} du}{\int_{S(v, \sigma)} e^{-\beta TV(u)} du}.$$

We conjecture that the three formulations of \hat{u}_{LSE} (Lagrangian, ball-constrained and sphere-constrained formulations) are asymptotically equivalent when the image dimensions are large enough. This conjecture comes from several remarks.

- Propositions 1.1 and 1.2 prove a certain form of equivalence between the normal distribution and the uniform distribution on the sphere, and between the uniform distribution on the sphere and the uniform distribution on the ball, respectively, in high dimension. Moreover Hida and Nomoto [1964] proved that the uniform distribution on the unit sphere of \mathbb{R}^n converges in distribution towards an infinite dimensional standard normal distribution when $n \rightarrow \infty$. It means that in high dimension, the uniform distribution on the unit sphere, the uniform distribution on the unit ball and the normal distribution are practically equivalent.

- It is true that for every positive sequence $(\varepsilon_n) \rightarrow 0$, the normal distribution $\mathcal{N}(v, \sigma^2 I_n)$ and the uniform distribution on $B(v, \sigma + \varepsilon_n)$ satisfy

$$d_{TV}(\mathcal{N}(v, \sigma^2 I_n), \mathcal{U}_{B(v, \sigma + \varepsilon_n)}) \xrightarrow{n \rightarrow \infty} 0$$

where $d_{TV}(\mu_1, \mu_2) = \sup_A |\mu_1(A) - \mu_2(A)|$ denotes the total variation distance between measures μ_1 and μ_2 . If the mappings $u \mapsto e^{-\beta TV(u)}$ and $u \mapsto ue^{-\beta TV(u)}$ were both bounded then we would get

$$\begin{cases} \frac{\int e^{-\beta TV(u)} e^{-\frac{\|u-v\|^2}{2\sigma^2}} du}{\int e^{-\frac{\|u-v\|^2}{2\sigma^2}} du} - \frac{\int_{B(v, \sigma + \varepsilon_n)} e^{-\beta TV(u)} du}{\int_{B(v, \sigma + \varepsilon_n)} du} \xrightarrow{n \rightarrow \infty} 0, \\ \frac{\int ue^{-\beta TV(u)} e^{-\frac{\|u-v\|^2}{2\sigma^2}} du}{\int e^{-\frac{\|u-v\|^2}{2\sigma^2}} du} - \frac{\int_{B(v, \sigma + \varepsilon_n)} ue^{-\beta TV(u)} du}{\int_{B(v, \sigma + \varepsilon_n)} du} \xrightarrow{n \rightarrow \infty} 0, \end{cases}$$

which would yield an equivalence between the Lagrangian \hat{u}_{LSE} and the ball-constrained formulation in $B(v, \sigma + \varepsilon_n)$ with $\varepsilon_n \rightarrow 0$. Now the mapping $u \mapsto ue^{-\beta TV(u)}$ is not bounded, and so this argument is not accurate.

In the sequel we shall keep the Lagrangian formulation of TV-LSE.

1.3.2 TV-LSE as a diffusion

Let us write \hat{u}_{LSE} into the following form

$$\hat{u}_{LSE} = \frac{1}{Z} \int ue^{-\frac{\|u-v\|^2}{2\sigma^2}} dp_\beta(u).$$

This is the convolution of the noisy image v with a Gaussian kernel. The convolution is carried out on the whole space of images, against the prior measure $dp_\beta(u)$, which favors rare but sharp edges. This can be considered as a diffusion of v with privileged features enhanced.

This writing is linked to the Non-Local Means algorithm proposed by [Buades et al. 2005]. Indeed, the Non-Local means can be interpreted as a diffusion in the space of patches, because it approximates the Laplace-Beltrami operator on this space equipped with the Euclidean distance [Lafon et al. 2006; Peyré 2008; Buades et al. 2007].

Chapter 6 will tackle the link between TV-LSE and NL-means more precisely.

1.3.3 Link to Wiener filter

The Wiener filter is also a denoising scheme based on mean square error minimization. It is then worth driving a comparison between TV-LSE and Wiener filter.

The Wiener filter assumes both the original image and the noise to be drawn from a weakly stationary stochastic process, i.e. a stationary processes whose autocovariance functions

$$c_U(t, s) = \text{cov}(U(t), U(s)) \quad \text{and} \quad c_\varepsilon(t, s) = \text{cov}(\varepsilon(t), \varepsilon(s)) \quad (t, s \in \Omega)$$

are functions of $s - t$ only. When the autocovariance functions c_U and c_ε are known, the Wiener filter $v \mapsto \hat{u}$ is defined as the linear and translation invariant operator which minimizes the mean square error

$$\int_{\Omega} (\hat{u}(t) - u(t))^2 dt.$$

When the noise is white and Gaussian, the Wiener filter can be written as a convolution by a bandpass kernel k depending on c_U and c_ε , namely

$$k(x, y) = \frac{P_U(x, y)}{P_U(x, y) + \sigma^2},$$

where P_U is the power spectrum of the image process, i.e. the Fourier transform of c_U .

If the model for the original image is the Total Variation distribution, then u is considered to be the realization of the Gibbs field $U \sim \frac{1}{Z} e^{-\beta TV(U)}$. But a Gibbs field is a special case of weakly stationary process [Grimmett and Stirzaker 2001]. Besides, the model for the noise ε is white and Gaussian, which is clearly a weakly stationary process. The autocovariance functions of U and ε allow to build a Wiener filter, which is a weaker denoiser than TV-LSE for several reasons. First, the Wiener filter, unlike TV-LSE, is linear. In particular the denoising of any image can be deduced from the denoising of a Dirac image. This means that the model behind it is poor. Wiener filter acts like a bandpass filter; this is not very well adapted to image purpose because it creates blur. Furthermore, knowing all the autocovariance functions is weaker than knowing the entire prior distribution. This means that the model underlying the Wiener filter is much poorer than TV-LSE. In Wiener filtering, the good properties of total variation are forgotten and the contours in images get blurry.

Chapter 2

MCMC algorithm for TV-LSE denoising

Résumé Dans ce chapitre un algorithme complet pour le calcul d'images débruitées par TV-LSE est détaillé. Il s'agit d'un cas particulier de l'algorithme de Hastings-Metropolis à marche aléatoire avec balayage aléatoire des pixels. La vitesse de convergence de cet algorithme est prévue lente. C'est pourquoi nous donnons un critère précis de convergence, basé sur une décomposition biais-variance de l'erreur quadratique moyenne. Nous discutons le réglage de chacun des paramètres intervenant dans la méthode, et donnons des algorithmes spécifiques à l'estimation de ces paramètres, ainsi que des justifications théoriques.

Abstract In this chapter a complete algorithm for the computation of TV-LSE denoised images is detailed. It is a particular case of the Hastings-Metropolis algorithm based on random walks with a random scan of the pixels. The convergence rate of this algorithm is expected to be quite slow. That is why an accurate convergence criterion is given, based on a bias-variance decomposition of the squared error expectation. The tuning of all the parameters of the method apart from the hyperparameters is discussed, and specific algorithms estimating each of these parameters are given along with theoretical background.

Let us consider a noisy image v again, defined on a discrete domain $\Omega \subset \mathbb{Z}^2$. As the real-life images are very big (up to several millions pixels), the computation of

$$\hat{u}_{LSE} = \frac{\int_{\mathbb{R}^{\Omega}} u e^{-\frac{\|u-v\|^2 + \lambda TV(u)}{2\sigma^2}} du}{\int_{\mathbb{R}^{\Omega}} e^{-\frac{\|u-v\|^2 + \lambda TV(u)}{2\sigma^2}} du} \quad (2.1)$$

requires integrations on high-dimensional spaces. We focus on methods able to approximate these integrals. Approximation by regular sampling is not feasible because an exhaustive summation over a grid in such a large dimension would require huge computation times. We would rather use the global regularity of our problem, in particular the fact that many images have negligible weight in the integral, which tells that a complete exploration of the image space is useless. Stochastic methods are able to tackle such high-dimensional integrals. In this chapter we focus on Markov Chain Monte-Carlo (MCMC) algorithms, which fit our problem particularly well.

2.1 MCMC and Hastings-Metropolis algorithms

We now recall some theoretical principles underlying the MCMC algorithms [Hastings 1970; Durrett 1996; Robert 1996; Grimmett and Stirzaker 2001]. A MCMC algorithm is any method producing an ergodic Markov chain with given stationary distribution π . The next two theorems explain the link to high-dimensional integrals.

Theorem 2.1 (Discrete ergodic theorem) [Grimmett and Stirzaker 2001; Ouvrard 2000] *Let (Υ_n) be an irreducible Markov chain defined on a countable set. If (Υ_n) has stationary distribution π , then for any function $g \in L^1(d\pi)$, i.e. such that*

$$\sum_u |g(u)|\pi(u) < \infty,$$

the convergence holds

$$\frac{1}{n} \sum_{k=1}^n g(\Upsilon_k) \xrightarrow[n \rightarrow \infty]{} \sum_u g(u)\pi(u) \quad a.s.$$

Hence a high-dimensional sum $\sum_u g(u)\pi(u)$ can be approximated by the Cesàro average $\frac{1}{n} \sum_{k=1}^n g(\Upsilon_k)$ of the sequence $(g(\Upsilon_n))$, where (Υ_n) is a Markov chain with stationary distribution π . Intuitively, the convergence holds because in its stationary behavior the chain visits each site u according to a frequency which is proportional to $\pi(u)$.

Theorem 2.1 concerns Markov chains defined on countable state spaces. It has also been generalized to Markov chains defined on uncountable state spaces. In this *continuous* framework too, the Cesàro average converges to a high-dimensional integral $\int g(u)d\pi(u)$. This is what we need for the computation of \hat{u}_{LSE} (Equation (2.1)). However as soon as the algorithm is implemented, the state space becomes discrete (and thus countable). That is why we will maintain both aspects (discrete/continuous setting), in order to justify our method both theoretically and practically.

The adaptation of Theorem 2.1 to the continuous setting (i.e. to the case of an uncountable state space) needs to generalize the concepts of irreducibility and recurrence to Markov chains defined on uncountable state spaces. Let us introduce ϕ -irreducibility and Harris-recurrence.

Definition 2.1 [Meyn and Tweedie 1993] *If ϕ is a measure, the chain (Υ_n) is said to be ϕ -irreducible if for any Borel set A ,*

$$\phi(A) > 0 \quad \Rightarrow \quad \forall x, \mathbb{P}(\exists n \geq 1, \Upsilon_n \in A \mid \Upsilon_0 = x) > 0.$$

(i.e. (Υ_n) visits A within finite time with non-null probability, whatever the initialization). More classically, if π is σ -finite, π is said to be a stationary measure for (Υ_n) (whose transition kernel is P) if

$$\forall A \text{ Borel set, } \pi(A) = \int P(u, A)d\pi(u).$$

Last, a chain (Υ_n) is said to be Harris-recurrent if (Υ_n) has a stationary measure π , (Υ_n) is π -irreducible and for each B with $\pi(B) > 0$

$$\forall u_0 \in \mathbb{R}^\Omega \quad \mathbb{P} \left(\sum_n \mathbf{1}_{\Upsilon_n \in B} = \infty \mid \Upsilon_0 = u_0 \right) = 1$$

(i.e. the chain visits B infinitely often, whatever the initialization).

Theorem 2.2 (Continuous ergodic theorem) [Meyn and Tweedie 1993] *If (Υ_n) is Harris-recurrent and admits a stationary measure π with total mass 1, then for any function $g \in L^1(d\pi)$, the convergence holds*

$$\frac{1}{n} \sum_{k=1}^n g(\Upsilon_k) \xrightarrow[n \rightarrow \infty]{} \int g(u) d\pi(u) \quad \text{a.s.}$$

In order to approximate the integral $\int g(u) d\pi(u)$, we need a method to simulate an irreducible Markov chain with given stationary distribution π . One way to build such a Markov Chain is known as Hastings-Metropolis algorithm, and described in the following propositions.

Proposition 2.1 (Discrete Hastings-Metropolis) [Grimmett and Stirzaker 2001] *Let π a discrete probability distribution, let $q(\cdot|\cdot)$ a conditional distribution (called increment distribution). Let (Υ_n) be the Markov chain defined by any initial measure whose support is inside the one of π , and the transition matrix*

$$P(u, u') = \begin{cases} \rho(u, u')q(u'|u) & \text{if } u' \neq u, \\ \sum_v (1 - \rho(u, v))q(v|u) & \text{if } u = u', \end{cases}$$

where

$$\rho(u, u') = \begin{cases} \min\left(\frac{\pi(u')q(u|u')}{\pi(u)q(u'|u)}, 1\right) & \text{if } \pi(u)q(u'|u) > 0, \\ 1 & \text{if } \pi(u)q(u'|u) = 0. \end{cases} \quad (2.2)$$

Then π is a stationary distribution for (Υ_n) .

Proof : The proof can be found in [Ouvrard 2000; Grimmett and Stirzaker 2001], but is simple and short enough to be written here. It suffices to prove that π is reversible with respect to P , i.e. that

$$\forall u, u', \quad \pi(u)P(u, u') = \pi(u')P(u', u). \quad (2.3)$$

Indeed, if π is reversible, then

$$\sum_u \pi(u)P(u, u') = \sum_u \pi(u')P(u', u) = \pi(u')$$

which proves that π is stationary.

Equation (2.3) clearly holds for $u = u'$. Now if $u \neq u'$ and $\pi(u')q(u|u') \leq \pi(u)q(u'|u)$, then

$$P(u, u') = q(u'|u) \cdot \frac{\pi(u')q(u|u')}{\pi(u)q(u'|u)},$$

and thus $\pi(u)P(u, u') = \pi(u')q(u|u')$ which equals $\pi(u')P(u', u)$ since $\pi(u)q(u'|u) \geq \pi(u')q(u|u')$. The same proof can be made if $\pi(u')q(u|u') \geq \pi(u)q(u'|u)$. Consequently, π is reversible with respect to P , thus stationary for the chain (Υ_n) . \square

In the continuous setting [Tierney 1994] provides the same kind of result. The distributions we need will all be measurable with respect to Lebesgue's measure. For convenience and in order to make the link to the discrete setting clearer, their density will be denoted by the same letter than the distribution. For instance π will denote both the probability distribution and its probability density function (p.d.f.).

Proposition 2.2 (Continuous Hastings-Metropolis) [Tierney 1994] *Let π a probability distribution. Let $q(\cdot|\cdot)$ a transition distribution. Let (Υ_n) be the Markov chain defined by any initial measure and the transition kernel*

$$P(u, u') = \begin{cases} \rho(u, u')q(u'|u) & \text{if } u' \neq u, \\ \int_v [1 - \rho(u, v)]q(v|u)dv & \text{if } u = u', \end{cases}$$

where ρ is defined as in Equation (2.2). Then π is a stationary distribution for (Υ_n) .

The proof is an adaptation of the one of the discrete case, and can be found in [Tierney 1994].

Both in the discrete and the continuous case, a transition $\Upsilon_k \rightarrow \Upsilon_{k+1}$ following kernel P is equivalent to an acceptance/rejection procedure:

$$\begin{aligned} & \text{draw } \Upsilon_{k+1/2} \sim q(\Upsilon_{k+1/2}|\Upsilon_k) \\ & \text{if } \pi(\Upsilon_{k+1/2})q(\Upsilon_k|\Upsilon_{k+1/2}) > \pi(\Upsilon_k)q(\Upsilon_{k+1/2}|\Upsilon_k) \quad \Upsilon_{k+1} = \Upsilon_{k+1/2} \\ & \text{otherwise } \Upsilon_{k+1} = \begin{cases} \Upsilon_{k+1/2} & \text{with probability } \frac{\pi(\Upsilon_{k+1/2})}{\pi(\Upsilon_k)} \frac{q(\Upsilon_k|\Upsilon_{k+1/2})}{q(\Upsilon_{k+1/2}|\Upsilon_k)}, \\ \Upsilon_k & \text{with probability } 1 - \frac{\pi(\Upsilon_{k+1/2})}{\pi(\Upsilon_k)} \frac{q(\Upsilon_k|\Upsilon_{k+1/2})}{q(\Upsilon_{k+1/2}|\Upsilon_k)}. \end{cases} \end{aligned}$$

To sum things up, if π is a discrete probability distribution, the approximation of a high-dimensional sum $\sum g(u)d\pi(u)$ by Hastings-Metropolis algorithm can be made by Algorithm 1, provided the chain is irreducible. In the continuous case, the approximation of $\int_u g(u)d\pi(u)$ can also be made by Algorithm 1, provided that the chain is Harris-recurrent.

To be efficient, this algorithm needs to satisfy several requirements:

- q , called the *increment distribution*, should be fast and easy to simulate,
- the case where q is symmetric (i.e. $\forall x, y, q(x|y) = q(y|x)$) is appreciated because it simplifies the computation of the ratio $\frac{\pi(\Upsilon_{k+1/2})}{\pi(\Upsilon_k)} \frac{q(\Upsilon_k|\Upsilon_{k+1/2})}{q(\Upsilon_{k+1/2}|\Upsilon_k)}$ into $\frac{\pi(\Upsilon_{k+1/2})}{\pi(\Upsilon_k)}$,
- n should be large enough for the returned value to be independent enough from the initial measure μ and for the Cesàro average to converge,
- it is not necessary to know the normalizing factor of π , since the Hastings-Metropolis algorithm only requires the computation of π through ratios $\frac{\pi(u')}{\pi(u)}$.

Algorithm 1 Hastings-Metropolis algorithm

choose an initial measure μ ,
 choose an increment distribution q easy to simulate,
 choose a maximal number of iterations n ,
 draw Υ_0 from initial measure μ ,
 set $k = 0$
repeat
 draw $\Upsilon_{k+1/2} \sim q(\Upsilon_{k+1/2}|\Upsilon_k)$,
 if $\pi(\Upsilon_{k+1/2})q(\Upsilon_k|\Upsilon_{k+1/2}) > \pi(\Upsilon_k)q(\Upsilon_{k+1/2}|\Upsilon_k)$ set $\Upsilon_{k+1} = \Upsilon_{k+1/2}$
 else set $\Upsilon_{k+1} = \begin{cases} \Upsilon_{k+1/2} & \text{with probability } \frac{\pi(\Upsilon_{k+1/2})}{\pi(\Upsilon_k)} \frac{q(\Upsilon_k|\Upsilon_{k+1/2})}{q(\Upsilon_{k+1/2}|\Upsilon_k)}, \\ \Upsilon_k & \text{with probability } 1 - \frac{\pi(\Upsilon_{k+1/2})}{\pi(\Upsilon_k)} \frac{q(\Upsilon_k|\Upsilon_{k+1/2})}{q(\Upsilon_{k+1/2}|\Upsilon_k)} \end{cases}$
 $k \leftarrow k + 1$
until $k = n + 1$
return $\frac{1}{n} \sum_{i=1}^n g(\Upsilon_i)$.

2.2 A Hastings-Metropolis algorithm for TV-LSE

Let us make the distributions π and $q(\cdot|\cdot)$ explicit in the continuous Hastings-Metropolis algorithm for the computation of \hat{u}_{LSE} given by Equation (2.1). This can also be done in the discrete setting; for details see the proof of Proposition 2.4.

- The Markov chain (Υ_n) is a sequence of images defined on Ω .
- The stationary density will be the posterior density

$$\pi(u) = \frac{1}{Z} e^{-\frac{\|u-v\|^2 + \lambda TV(u)}{2\sigma^2}}, \quad (2.4)$$

which is known to a constant multiple. In this section the operator TV can be very general, its only restriction is that π should be well defined, i.e. positive and integrable.

- The coordinate functions $g_x : u \mapsto u(x)$ defined for each pixel $x \in \Omega$ are all π -integrable (see Proposition 1.5). If (Υ_n) is Harris-recurrent, then the gray level $\hat{u}_{LSE}(x)$ can then be approximated by $\frac{1}{n} \sum_{k=1}^n g_x(\Upsilon_k)$.
- The density of the increment distribution $q(u'|u)$ will be chosen as follows. Let $\alpha > 0$, then $q(\cdot|\cdot)$ is defined by

$$q(u'|u) = \frac{1}{|\Omega|} \sum_{x \in \Omega} \left(\prod_{y \neq x} \delta_{u(y)}(u'(y)) \right) \frac{1}{2\alpha} \mathbf{1}_{[u(x)-\alpha, u(x)+\alpha]}(u'(x)) \quad (2.5)$$

It is equivalent to drawing a pixel $x \sim \mathcal{U}_\Omega$ uniformly on the set of pixels, and to allow a uniform random walk $\mathcal{U}_{[u(x)-\alpha, u(x)+\alpha]}$ on this pixel alone. This is a simple way to have q both symmetric and fast to simulate. Notice that $q(u'|u)$ is a function of $u' - u$, so

that this algorithm can be classified into the *random walk*-based Hastings-Metropolis algorithms. As the transition kernel allows one pixel at most to be modified, and as this pixel is uniformly drawn from Ω , this algorithm enters the category of the *variable-at-a-time* Hastings-Metropolis with *random scan* [Fort et al. 2003; Roberts and Rosenthal 1997, 1998].

The procedure is summed up in Algorithm 2.

Algorithm 2 Continuous Hastings-Metropolis algorithm for \hat{u}_{LSE}

let $\pi(u) = \exp[-\frac{\|u-v\|^2 + \lambda TV(u)}{2\sigma^2}]$
 choose a maximal number of iterations n ,
 choose a subsampling rate R ,
 draw Υ_0 at random from an initial distribution μ ,
 set $k = 0$
repeat
 draw $x \sim \mathcal{U}_\Omega$
 draw $\begin{cases} \Upsilon_{k+1/2}(x) \sim \mathcal{U}_{[\Upsilon_k(x)-\alpha, \Upsilon_k(x)+\alpha]} \\ \Upsilon_{k+1/2}(y) = \Upsilon_k(y) \quad \forall y \neq x \end{cases}$
 if $\pi(\Upsilon_{k+1/2}) > \pi(\Upsilon_k)$ set $\Upsilon_{k+1} = \Upsilon_{k+1/2}$
 else set $\Upsilon_{k+1} = \begin{cases} \Upsilon_{k+1/2} & \text{with probability } \frac{\pi(\Upsilon_{k+1/2})}{\pi(\Upsilon_k)} \\ \Upsilon_k & \text{with probability } 1 - \frac{\pi(\Upsilon_{k+1/2})}{\pi(\Upsilon_k)} \end{cases}$
 $k \leftarrow k + 1$
until $k = Rn + 1$
return $\frac{1}{n} \sum_{i=1}^n \Upsilon_{Ri}$.

The determination of good values for parameter n in Algorithm 2 will be the object of Section 2.3, where a precise and efficient stopping criterion will be given. The selection of parameter R has smaller impact, and is discussed now.

2.2.1 Subsampled chain

As two successive images Υ_n and Υ_{n+1} differ on one pixel at most, the evolution of the chain (Υ_n) is very slow. In the following we will consider the chain subsampled by a factor R , that is (Υ_{Rn}). This subsampled chain presents the following advantages:

- When R is large the correlation between the iterations is much smaller and may allow a faster convergence. Of course, using the subsampled chain requires the computation of all the intermediate states $(\Upsilon_{Rn+k})_{1 \leq k \leq R-1}$. Therefore the gain in practical convergence rate cannot be that meaningful.
- When $R = |\Omega|$ ($|\Omega|$ is the number of pixels in the datum image), the convergence rate of the subsampled chain is “normalized” by the number of pixels $|\Omega|$.

- The subsampled chain (Υ_{Rn}) converges to \hat{u}_{LSE} the same way as the complete chain (Υ_n) , for any $R \in \mathbb{N}^*$ (see Proposition 2.4 below).

From now on we shall call “iteration” the run of R loops in Algorithm 2. When $R = |\Omega|$, this corresponds to a random scan of $|\Omega|$ pixels, i.e. to an approximate image scan. In practice the chain has got the same convergence rate for a large range of subsampling parameters $R > 0$. The main advantage then is to have this “normalization” with respect to $|\Omega|$.

2.2.2 Consistency of Algorithm 2

Let us first recall a proposition from [Tierney 1994] which will simplify the proof of Proposition 2.4.

Proposition 2.3 [Tierney 1994, Corollary 2] *Suppose P is a π -irreducible Hastings-Metropolis kernel. Then P is Harris-recurrent.*

Now we come to the main result of this section.

Proposition 2.4 (Algorithm 2 converges) *Let (Υ_n) be the Markov chain constructed by Algorithm 2. Then for any $R \in \mathbb{N}^*$ its subsampled chain (Υ_{Rn}) satisfies*

$$\frac{1}{n} \sum_{k=1}^n \Upsilon_{Rk} \xrightarrow[n \rightarrow \infty]{} \hat{u}_{LSE} \quad a.s.$$

both in the discrete and the continuous cases.

Proof : Discrete framework: Let l be a positive real number (standing for a quantization step). The increment distribution is the discrete version of Eq. (2.5), that is

$$q(u'|u) = \frac{1}{|\Omega|} \sum_{x \in \Omega} \left(\prod_{y \neq x} \delta_{u(y)}(u'(y)) \right) \frac{1}{(2l+1)} \sum_{-l \leq k \leq l} \delta_{u(x) + \alpha k/l}(u'(x)),$$

where $\alpha > 0$ is the increment parameter of the chain. The chain (Υ_n) defined by Hastings-Metropolis algorithm associated to increment distribution q is irreducible. Indeed if u and u' are two images in $E = (\alpha\mathbb{Z}/l)^\Omega$, then

$$\forall n \geq \frac{|\Omega|}{\alpha} \cdot \|u - u'\|_\infty, \quad \mathbb{P}(\Upsilon_n = u' | \Upsilon_0 = u) > 0.$$

In particular, if $n \geq \left\lceil \frac{|\Omega|}{\alpha R} \cdot \|u - u'\|_\infty \right\rceil$, then $\mathbb{P}(\Upsilon_{Rn} = u' | \Upsilon_0 = u) > 0$, so that u and u' communicate in the subsampled chain.

The chain (Υ_n) has stationary distribution

$$\pi(u) = \frac{1}{Z} e^{-\frac{\|u-v\|^2 + \lambda TV(u)}{2\sigma^2}}$$




Figure 2.1: Construction of the sequence (u_1, \dots, u_m) in the proof of the π -irreducibility of (Υ_n) in the continuous framework. We start from an initializing image u_0 and a Borel set A containing an open ball $B(u, \varepsilon)$. We then build a sequence of balls beginning from $B(u_0, \varepsilon)$ until $B(u, \varepsilon)$ whose successive intersections are non-empty, and for which the centers of two successive balls differ by only one pixel at most.

defined on E because its transition kernel is reversible with respect to π (Proposition 2.1). π is also stationary for (Υ_{Rn}) , because (Υ_{Rn}) has transition matrix P^R , and $\pi P = \pi$ implies $\pi P^R = \pi$.

Then thanks to Theorem 2.1 the first convergence holds.

Continuous framework: Let us prove that (Υ_n) is π -irreducible. Let us first introduce

$$c_P(B_1|B_0) = \inf_{u \in B_0} \int_{B_1} P(u'|u) du'$$

for B_1 any Borel subset, and B_0 any subset of $E = \mathbb{R}^\Omega$. The quantity $c_P(B_1|B_0)$ represents the minimum connectivity from B_0 to B_1 for transition P .

Let A be a Borel subset of $E = \mathbb{R}^\Omega$, such that $\pi(A) > 0$. Then A contains an open ball $B(u, \varepsilon)$ of radius ε and centered at u ; we can assume $\varepsilon \leq \alpha/4$ (α is the increment parameter of the chain, see Eq. (2.5)). Let u_0 be an initial image. As E is connected, there exists a sequence $(u_1, \dots, u_m, \dots, u_{m+R})$ of images such that

$$\begin{cases} \forall k \in \{0, \dots, m-1\}, u_k \text{ and } u_{k+1} \text{ differ on 1 pixel at most, denoted } x_k \\ \forall k \in \{0, \dots, m-1\}, |u_{k+1}(x_k) - u_k(x_k)| \leq 2\varepsilon \\ u_m = u_{m+1} = \dots = u_{m+R} = u \in A \end{cases}$$

(see Figure 2.1). Then each step $B(u_k, \varepsilon) \rightarrow B(u_{k+1}, \varepsilon)$ has a non-zero probability for the increment distribution, i.e.

$$\forall k \in \{0, \dots, m+R-1\}, \inf_{u \in B(u_k, \varepsilon)} \int_{B(u_{k+1}, \varepsilon)} q(u'|u) du' > 0.$$

Furthermore, as $\pi(B(u_{k+1}, \varepsilon)) > 0$ and $\pi(B(u_k, \varepsilon)) > 0$, all the transitions have a positive probability to be accepted under P , i.e.

$$\forall k \in \{0, \dots, m+R-1\}, c_P(B(u_{k+1}, \varepsilon) | B(u_k, \varepsilon)) > 0.$$

Finally the probability that the chain (Υ_{Rn}) initiated at u_0 will be visiting A in finite time satisfies

$$\begin{aligned} \mathbb{P}(\exists n \geq 1, \Upsilon_{Rn} \in A | \Upsilon_0 = u_0) &\geq \mathbb{P}(\Upsilon_{R\lfloor \frac{m+R}{R} \rfloor} \in A | \Upsilon_0 = u_0) \geq \mathbb{P}(\Upsilon_{R\lfloor \frac{m+R}{R} \rfloor} \in B(u, \varepsilon) | \Upsilon_0 = u_0) \\ &\geq \int_{B(u_m, \varepsilon)} P(\theta_m | \theta_{m-1}) d\theta_m \int_{B(u_{m-1}, \varepsilon)} P(\theta_{m-1} | \theta_{m-2}) d\theta_{m-1} \cdots \int_{B(u_1, \varepsilon)} P(\theta_1 | u_0) d\theta_1 \\ &\geq \left[\prod_{k=1}^{\lfloor \frac{m+R}{R} \rfloor - 1} c_P(B(u_{k+1}, \varepsilon) | B(u_k, \varepsilon)) \right] \cdot c_P(B(u_1, \varepsilon) | \{u_0\}) > 0. \end{aligned}$$

Then (Υ_{Rn}) is π -irreducible.

Besides π is stationary for (Υ_n) because of Proposition 2.2.

Now thanks to Proposition 2.3 the subsampled chain is Harris-recurrent. Consequently, thanks to Theorem 2.2 the convergence holds. \square

The random walk-based Hastings-Metropolis method described in Algorithm 2 is now proven to converge towards \hat{u}_{LSE} . The method has several parameters whose tuning can be crucial for the convergence speed. Namely,

- the subsampling factor R will be chosen as $R = |\Omega|$ in practice, so that the convergence rate of the chain studied in Section 2.3 is normalized by the number of pixels $|\Omega|$. In the sequel

$$U_n = \Upsilon_{|\Omega|n} \tag{2.6}$$

will denote the subsampled chain,

- the total number of iterations n required for the convergence of the process will be treated in Section 2.3, relative to the convergence control of the process,
- the choice of the initial distribution μ will be discussed in Section 2.4 through the viewpoint of burn-in of the chain,
- the increment distribution depends on parameter α , whose tuning is discussed in Section 2.5.

2.3 Convergence control

This section is devoted to the construction of an accurate convergence criterion of the chain. The results given are rather technical and rely on recent Statistics papers. Let us briefly summarize the section. First, we give arguments in favor of a rather slow convergence of the chain, which comes from a generalization of a Central Limit Theorem. This motivates us to build an efficient convergence criterion. Then the approximation error made by the chain is decomposed through bias and variance terms. Both terms are carefully bounded thanks to state-of-the-art results of Statistics. Last, we introduce a second chain, independent from the first one and identically

distributed (i.i.d.). A comparison between the two chains is then used to make an efficient stopping criterion.

In this section we only maintain the continuous point of view. A numerical scheme for total variation (TV) does not need to be specified exactly, but the operator TV is assumed to satisfy the following requirements.

Requirements for TV operator: If Ω is a finite subset of \mathbb{Z}^2 , a TV operator is an operator $\mathbb{R}^\Omega \rightarrow \mathbb{R}$ such that

1. the posterior distribution $\pi(u) = \frac{1}{Z} e^{-\frac{\|u-v\|^2 + \lambda TV(u)}{2\sigma^2}}$ is absolutely continuous with respect to Lebesgue's measure on \mathbb{R}^Ω , with positive and continuous density on \mathbb{R}^Ω . This regularity condition is quite natural to fit the classical framework for MCMC algorithms, and will be useful for Lemma 2.1.
2. the posterior distribution π admits finite second order moments. This will be useful to measure the speed of the MCMC algorithm, thanks to a Central Limit Theorem (Theorem 2.3).
3. TV is Lipschitz continuous, i.e. there exists $\beta > 0$ such that

$$\forall u \in \mathbb{R}^\Omega, \forall h \in \mathbb{R}^\Omega, \quad |TV(u+h) - TV(u)| \leq \beta \|h\| \quad (2.7)$$

(for Lemma 2.2).

4. for any $s \in (0, 1)$, π^{-s} is “superlinear”, i.e. there exists $a > 0$ such that

$$\forall u \in \mathbb{R}^\Omega, \quad \pi^{-s}(u) \geq a \|u\|. \quad (2.8)$$

This condition forbids $\pi(u)$ to converge to 0 too slowly when $\|u\| \rightarrow \infty$, and will be required for Proposition 2.5.

The most common numerical schemes for TV , such as

$$\sum_{x,y} (|u(x+1,y) - u(x,y)| + |u(x,y+1) - u(x,y)|)$$

or

$$\sum_{x,y} \sqrt{|u(x+1,y) - u(x,y)|^2 + |u(x,y+1) - u(x,y)|^2},$$

satisfy these requirements. More generally, the requirements are satisfied as soon as TV is convex and Lipschitz continuous (2.7) (see Chapter 3, Section 3.4 for more details).

2.3.1 Central limit theorem

Now that the convergence of the subsampled chain average $\frac{1}{n} \sum_{k=1}^n U_k$ towards \hat{u}_{LSE} has been proven, one can wonder about the convergence rate of such a method. If the random variables (U_n) were independent and identically distributed, then the Central Limit Theorem could apply and give a convergence rate in $1/\sqrt{n}$. As here (U_n) is a Markov chain, we need a generalization of the Central Limit Theorem to non i.i.d. random variables to check if the rate is in $1/\sqrt{n}$.

Kipnis and Varadhan [1986]; Roberts and Tweedie [1996] provide this kind of result. Before calling back one of their results, let us give the following definition.

Definition 2.2 *The Markov chain (U_n) is said to satisfy the Central Limit Theorem for $g : \mathbb{R}^n \rightarrow \mathbb{R}$, if there exists γ_g such that*

$$\frac{1}{\sqrt{n}} \left(\sum_{k=1}^n \left[g(U_k) - \int g(u) d\pi(u) \right] \right) \xrightarrow[n \rightarrow \infty]{\mathcal{D}} \mathcal{N}(0, \gamma_g^2).$$

($\xrightarrow{\mathcal{D}}$ means a convergence in distribution).

Kipnis and Varadhan [1986] state a Central Limit Theorem in the case where the chain is stationary (i.e. initialized by $U_0 \sim \pi$) and reversible. More precisely,

Theorem 2.3 [Kipnis and Varadhan 1986] *Assume that the chain (U_n) is stationary (i.e. initialized by $U_0 \sim \pi$), π -irreducible and reversible. Let g be a square integrable with respect to π , and let γ_g be defined by*

$$\begin{aligned} \gamma_g^2 &= \lim_{n \rightarrow \infty} n \mathbb{E} \left[\left(\frac{1}{n} \sum_{k=1}^n \bar{g}(U_k) \right)^2 \right] \\ &= \mathbb{E}[\bar{g}^2(U_0)] + 2 \sum_{k=1}^{+\infty} \mathbb{E}[\bar{g}(U_0) \bar{g}(U_k)], \end{aligned} \tag{2.9}$$

where \bar{g} is the centered version of g , i.e. $\bar{g} = g - \int g(u) d\pi(u)$. Then when $0 < \gamma_g < \infty$, the chain (U_n) satisfies the Central Limit Theorem for g .

The application of Theorem 2.3 is not straightforward:

- The theorem is stated in the case where the chain is initialized from π . Now, [Geyer 1992, Remark 1.1] tells that if the chain is Harris recurrent, the convergence does not depend on the starting measure of the chain. As in our case the chain is Harris recurrent, the chain need not be stationary to satisfy Theorem 2.3.
- In our case, the considered functions g will be the coordinate functions $g_x : u \mapsto u(x)$ for each pixel x . The coordinate functions $(g_x)_{x \in \Omega}$ are square integrable with respect to π .
- One needs to check whether $0 < \gamma_g < \infty$ or not. This is a tough problem which in practice requires to estimate γ_g [Geyer 1992]. In the following we will assume that $0 < \gamma_g < \infty$ for every coordinate function $g_x : u \mapsto u(x)$ ($x \in \Omega$).

Under this last assumption, we can apply Theorem 2.3, and the convergence

$$\frac{1}{n} \sum_{k=1}^n U_k \xrightarrow[n \rightarrow \infty]{} \hat{u}_{LSE}$$

occurs at rate γ/\sqrt{n} for the normalized L^2 -norm, where

$$\gamma^2 = \frac{1}{|\Omega|} \sum_{x \in \Omega} \gamma_{g_x}^2.$$

Let us introduce the notion of *span*, which we shall need in the sequel. It is a very natural generalization of the variance for random vectors.

Definition 2.3 Let U be a random vector defined on \mathbb{R}^n . Its span is defined by

$$\text{Span}(U) = \text{Tr}(\text{Cov}(U)),$$

i.e. by the trace of its covariance matrix.

Remark: Both ergodic theorem and central limit theorem hold for any initializing distribution, upon condition that γ is defined on the stationary regime of (U_n)

$$\gamma = \lim_{b \rightarrow \infty} \lim_{n \rightarrow \infty} (n - b) \text{Span} \left(\frac{1}{n - b} \sum_{k=b+1}^n U_k \right).$$

However in practice the convergence rate highly depends on this initialization. Indeed, as π is continuous and log-concave, if the chain is badly initialized (*i.e.* $\pi(U_0)$ is very low), then the random walk-based Hastings-Metropolis algorithm will take many iterations for the chain to get into regions where π is high enough. All these iterations will be in the Cesàro average and will pollute the estimation for a very long time before they become negligible. Now this problem can be overcome by forgetting the first iterations of the chain in the average, as described in the next paragraph.

Note that parameters γ close to zero yield faster convergence large γ . The subsampled chain (U_n) (Eq. (2.6)) is expected to have correlations $\mathbb{E}_\pi[\bar{g}_x(U_0)\bar{g}_x(U_k)]$ closer to zero than the whole chain (U_n) . It should then converge faster with respect to the number n of large iterations (each containing R runs in the loop of Algorithm 2), but not necessarily faster when considering the total number of runs Rn .

2.3.2 Bias-variance decomposition

A convergence rate in $1/\sqrt{n}$ is quite slow; moreover dealing with large images (*e.g.* 512×512 , which is very small yet for real-life images) makes an algorithm work very slowly in practice. That is why an accurate convergence control is needed.

Let

$$S_n^b = \frac{1}{n - b} \sum_{k=b+1}^n U_k \tag{2.10}$$

be the Cesàro average of the subsampled chain where the b first terms have been deleted. The parameter b is a *burn-in* parameter, and allows the chain to take the time to reach more or less the stationary regime before the averaging process begins. It is classically used by MCMC practitioners because it reduces the impact of the first values of the chain when the stationary distribution π is not easy to simulate (see remark above). Taking b too small runs the risk of slowing down the practical convergence rate, when the chain is not properly initialized. Conversely, taking b too large reduces the number of terms in the Cesàro average and makes the approximation less precise. So there must exist a middle b attaining optimality. Such a trade-off can be explained as a bias-variance trade-off.

Indeed the expected distance from our estimate S_n^b to the target image \hat{u}_{LSE} satisfies the following bias-variance decomposition

$$\begin{aligned}\mathbb{E}\|S_n^b - \hat{u}_{LSE}\|^2 &= \mathbb{E}\|S_n^b - \mathbb{E}S_n^b\|^2 + \mathbb{E}\|\mathbb{E}S_n^b - \hat{u}_{LSE}\|^2 + 2\mathbb{E}\langle S_n^b - \mathbb{E}S_n^b, \mathbb{E}S_n^b - \hat{u}_{LSE} \rangle \\ &= \mathbb{E}\|S_n^b - \mathbb{E}S_n^b\|^2 + \|\mathbb{E}S_n^b - \hat{u}_{LSE}\|^2,\end{aligned}\tag{2.11}$$

where the first term $\mathbb{E}\|S_n^b - \mathbb{E}S_n^b\|^2$ is the span of S_n^b , and the rightmost term $\|\mathbb{E}S_n^b - \hat{u}_{LSE}\|^2$ stands for its squared bias. Let us derive both terms separately.

First a more precise central limit theorem applied to (S_n^b) would allow to control the span term. Indeed a central limit theorem would state a convergence

$$\sqrt{n-b} \left[S_n^b - \mathbb{E}S_n^b \right] \xrightarrow[n-b \rightarrow \infty]{\mathcal{D}} \mathcal{N}(0, \text{diag}(\gamma_x^2)),$$

where here the variance coefficients (γ_x) would be defined on the subsampled chain by

$$\gamma_x^2 = \lim_{n \rightarrow \infty, n-b \rightarrow \infty} n \mathbb{E}[S_n^b(x) - \hat{u}_{LSE}(x)]^2.$$

Unfortunately, to our knowledge no statistical result was brought about it. In the sequel we assume that it is true.

Assuming the images S_n^b to have large dimensions and to be drawn from a stationary process, we have $\|S_n^b - \mathbb{E}S_n^b\|^2 \approx \mathbb{E}\|S_n^b - \mathbb{E}S_n^b\|^2 = \text{Span}(S_n^b)$ thanks to Bienaymé-Chebyshev inequality for instance. Now with high probability the span term can be bounded from above for a certain $A > 0$ by

$$\text{Span}(S_n^b) \leq \frac{A}{n-b}.\tag{2.12}$$

Let us now concentrate on the bias term $\|\mathbb{E}S_n^b - \hat{u}_{LSE}\|$. We can prove that the chain (Υ_n) is geometrically ergodic, i.e. that there exists a function $V \geq 1$ a.e. such that $\pi(V) < \infty$, and there exist constants $r > 1$ and $R < \infty$ such that for π -almost all u ,

$$\forall n \in \mathbb{N}, \|P^n(u, \cdot) - \pi\|_V \leq RV(u)r^{-n},$$

where, P denotes the transition kernel of (Υ_n) , and, for any signed measure μ ,

$$\|\mu\|_V := \sup_{|f| \leq V} |\mu(f)|.$$

To proceed, we apply the result from [Fort et al. 2003] which gives conditions to the chain to be geometrically ergodic. These conditions fit our framework particularly well. Indeed, there is no requirement for π to be regular, contrary to most results [Jarner and Hansen 2000; Roberts and Tweedie 1996; Roberts and Rosenthal 1997, 1998] which require differentiability of π and conditions on its curvature. We recall this result as a lemma.

Lemma 2.1 [Fort et al. 2003, Theorem 2] *Let (Υ_n) be a symmetric random walk-based Hastings-Metropolis chain defined on \mathbb{R}^d , where the increment distribution is a “variable-at-a-time” distribution with “random scan” (the pixels are visited through a random uniform order). Three conditions need to be satisfied:*

1. Assume that the target distribution π is absolutely continuous with respect to Lebesgue's measure on \mathbb{R}^d , with positive and continuous density on \mathbb{R}^d .
2. Let $\{q_x\}_{1 \leq x \leq d}$ be the increment distribution family of symmetric densities absolutely continuous with respect to Lebesgue's measure on \mathbb{R} (q_x denotes the increment distribution related to pixel x). Assume that there exist constants η_x and $\delta_x < \infty$ ($1 \leq x \leq d$) with the property that $q_x(y) \geq \eta_x$ whenever $|y| \leq \delta_x$.
3. Assume that there are δ and Δ with $0 \leq \delta < \Delta \leq +\infty$ such that

$$\xi := \inf_{1 \leq x \leq d} \int_{\delta}^{\Delta} q_x(y) dy > 0,$$

and there exist functions $\phi_k \in \Phi$ for $1 \leq k \leq d$, such that, for all $1 \leq x \leq d$ and all $y \in [\delta, \Delta]$,

$$\begin{cases} \lim_{|u_x| \rightarrow \infty} \sup_{\phi_z(|u_z|) \leq \phi_x(|u_x|), z \neq y} \frac{\pi(u)}{\pi(u - \text{sign}(u_x)ye_x)} = 0 \\ \lim_{|u_x| \rightarrow \infty} \sup_{\phi_z(|u_z|) \leq \phi_x(|u_x|), z \neq y} \frac{\pi(u + \text{sign}(u_x)ye_x)}{\pi(u)} = 0, \end{cases}$$

with e_x the coordinate unit vector associated to the x -th pixel, and Φ defined as the set of functions $\phi : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ such that

$$\lim_{n \rightarrow \infty} \phi(t_n) = +\infty \quad \text{if and only if} \quad t_n \rightarrow +\infty.$$

Let $r : (0, 1) \rightarrow \mathbb{R}^+$ defined as

$$r(s) = 1 + s(1 - s)^{1/s-1}.$$

Let $s \in (0, 1)$ such that

$$r(s) < 1 + \frac{\xi}{d - 2\xi}, \tag{2.13}$$

and set

$$V_s(u) = \pi(u)^{-s}. \tag{2.14}$$

Then the chain is V_s -geometrically ergodic, i.e. there exists $R < \infty$ such that for any initialization u ,

$$\forall n \in \mathbb{N}^*, \quad \|P^n(u, \cdot) - \pi\|_{V_s} \leq Rr(s)^{-n}V_s(u).$$

Let us apply this lemma to our TV framework.

Lemma 2.2 *There exist values of s for which the assumptions of Lemma 2.1 are satisfied by our chain (Y_n) , and then for which (Y_n) is V_s -geometrically ergodic.*

Proof : Let us check the 3 items of the assumptions of Lemma 2.1.

1. $\pi = \frac{1}{Z} e^{-\frac{E_\lambda}{2\sigma^2}}$ is clearly a positive and continuous density on \mathbb{R}^d with $d = |\Omega|$.

2. Here the q_x are all equal to $\frac{1}{2\alpha}\mathbf{1}_{[-\alpha,\alpha]}$. Then taking $\delta_x = \alpha$ and $\eta_x = \frac{1}{2\alpha}$ validates the second assumption.
3. Let $\delta = \alpha/2$ and $\Delta = +\infty$; then $\xi = 1/4$ which is positive. Afterwards, let (ϕ_x) all equal to identity. Then for any $1 \leq x \leq d$ and any $y \geq \alpha/2 > 0$, thanks to the Lipschitz continuity of operator TV (2.7),

$$\begin{aligned} \frac{\pi(u)}{\pi(u - \text{sign}(u_x)ye_x)} &= \exp\left(-\frac{E_\lambda(u) - E_\lambda(u - \text{sign}(u_x)ye_x)}{2\sigma^2}\right) \\ &\leq \exp\left(-\frac{2\text{sign}(u_x)y(u_x - v_x) - y^2 - \lambda\beta\|e_x\||y|}{2\sigma^2}\right) \xrightarrow{|u_x| \rightarrow \infty} 0 \end{aligned}$$

(β is the Lipschitz parameter of TV) and similarly

$$\begin{aligned} \frac{\pi(u + \text{sign}(u_x)ye_x)}{\pi(u)} &= \exp\left(-\frac{E_\lambda(u + \text{sign}(u_x)ye_x) - E_\lambda(u)}{2\sigma^2}\right) \\ &\leq \exp\left(-\frac{2\text{sign}(u_x)y(u_x - v_x) + y^2 - \lambda\beta\|e_x\||y|}{2\sigma^2}\right) \xrightarrow{|u_x| \rightarrow \infty} 0. \end{aligned}$$

Then the third condition is satisfied.

Function r is increasing from 1 to 2, and $1 + \frac{\xi}{d-2\xi} = 1 + \frac{1}{4d-2}$ is larger than but close to 1 (d is large). Then condition (2.13) is satisfied for any

$$s \in (0, r^{-1}(1 + \frac{1}{4d-2})). \quad (2.15)$$

As V_s is π -integrable for any $s \in (0, 1)$, the chain is then geometrically ergodic with $V_s = \pi^{-s}$ as soon as (2.15) is satisfied. \square

The chain is then geometrically ergodic, which allows to control the convergence speed of the chain expectations $(\mathbb{E}U_n)$ towards \hat{u}_{LSE} , as shows the following proposition. This will provide a bound for the bias term in the bias-variance decomposition.

Proposition 2.5 *Let s as in Eq. (2.15) (where $d = |\Omega|$). Suppose that the chain (U_n) is initialized by an initial distribution μ_0 such that $\mu_0(\pi^{-s}) < \infty$. Then there exist $\delta \in (0, 1)$ and $B > 0$ such that*

$$\forall n \in \mathbb{N}, \quad \|\mathbb{E}U_n - \hat{u}_{LSE}\| \leq B\delta^{|\Omega|n}.$$

Proof : Let s as in (2.15) (where $d = |\Omega|$). Then thanks to Lemma 2.2, (Y_n) is V_s -geometrically ergodic with $V_s = \pi^{-s}$. Let $(g_x)_{x \in \Omega}$ be the coordinate functions

$$g_x : u \mapsto u(x).$$

As $s > 0$, V_s is ‘‘superlinear’’ (2.8), and there exists a independent from $x \in \Omega$ such that

$$\forall u \in \mathbb{R}^\Omega, \quad |g_x(u)| \leq aV_s(u).$$

Then for any $n \in \mathbb{N}^*$, if the chain is initialized by distribution μ_0 , we have

$$\begin{aligned} \|\mathbb{E}\Upsilon_n - \hat{u}_{LSE}\|^2 &= \sum_{x \in \Omega} \left| \int P^n(u_0, g_x) d\mu_0(u_0) - \hat{u}_{LSE}(x) \right|^2 \\ &\leq \sum_{x \in \Omega} \left(\int |P^n(u_0, g_x) - \hat{u}_{LSE}(x)| d\mu_0(u_0) \right)^2 \\ &\leq \sum_{x \in \Omega} \left(\int aRr(s)^{-n} V_s(u_0) d\mu_0(u_0) \right)^2 \end{aligned}$$

because the chain is geometrically ergodic, and $|g_x/a| \leq V_s$. Now letting

$$B = \sqrt{|\Omega|} aR \cdot \int V_s d\mu_0 < \infty \quad \text{and} \quad \delta = r(s)^{-1} \in (0, 1),$$

we get

$$\|\mathbb{E}\Upsilon_n - \hat{u}_{LSE}\| \leq B\delta^n,$$

and letting $n = |\Omega|n'$, we get the desired result. \square

Remark: In Proposition 2.5, the assumption about the initial distribution requiring that

$$\mu_0(\pi^{-s}) < \infty$$

is true in practical cases. Indeed in our experiments we use either deterministic initializations ($u_0 = v$ or $u_0 = \hat{u}_{MAP}$) for which $\mu_0(\pi^{-s}) = \pi^{-s}(u_0) < \infty$, or pure noise initializations, where the gray levels are independent and uniformly drawn from interval $[0, 255]$. Then μ_0 has bounded support, and as $\pi^{-s}(u_0)$ is everywhere finite, $\mu_0(\pi^{-s}) < \infty$ holds. However if μ_0 is a normal distribution $\mathcal{N}(u', \tau^2 Id)$, then the condition holds iff $\tau^2 < \sigma^2/s$. We need to remain watchful in the choice of the initializing measure.

From the inequality of Proposition 2.5 we can deduce an upper bound for the bias term thanks to the triangle inequality

$$\begin{aligned} \left\| \mathbb{E}S_n^b - \hat{u}_{LSE} \right\| &= \left\| \frac{1}{n-b} \sum_{k=b+1}^n (\mathbb{E}U_k - \hat{u}_{LSE}) \right\| \\ &\leq \frac{1}{n-b} \sum_{k=b+1}^n \|\mathbb{E}U_k - \hat{u}_{LSE}\| \\ &\leq \frac{1}{n-b} \sum_{k=b+1}^n B\delta^{|\Omega|k} = \frac{B}{n-b} \delta^{|\Omega|(b+1)} \frac{1 - \delta^{|\Omega|(n-b)}}{1 - \delta^{|\Omega|}}. \end{aligned}$$

Denoting $B' = \left(\frac{B\delta^{|\Omega|}}{1 - \delta^{|\Omega|}} \right)^2$, we obtain the following upper bound for the squared bias

$$\left\| \mathbb{E}S_n^b - \hat{u}_{LSE} \right\|^2 \leq B' \frac{\delta^{2|\Omega|b}}{(n-b)^2}. \quad (2.16)$$

Last, thanks to (2.11) and (2.16), we obtain the main result of this subsection.

Theorem 2.4 *The expected distance from our estimate S_n^b to \hat{u}_{LSE} in the bias-variance decomposition is bounded by*

$$\mathbb{E}\|S_n^b - \hat{u}_{LSE}\|^2 \leq \text{Span } S_n^b + B' \frac{\delta^{2|\Omega|b}}{(n-b)^2}.$$

Assuming (2.12), we conclude that

$$\mathbb{E}\|S_n^b - \hat{u}_{LSE}\|^2 \leq \frac{A}{n-b} + B' \frac{\delta^{2|\Omega|b}}{(n-b)^2}. \quad (2.17)$$

We can note the following facts:

- The convergence rate when $b \rightarrow \infty$ and $n-b \rightarrow \infty$ is in $\sqrt{n-b}$, due to the variance term. We find back the Central limit theorem stated at the beginning of this section.
- For large n , the variance term prevails, while for small n , the bias can be dominating.
- For a given n , there exists a value of b in $[0, n)$ minimizing the error. The optimal value of b can be either 0 or a non-trivial value, depending on the values of A , B and δ . The optimal b may be very large, and may even exceed $n/2$, which means that more than half of the iterations may pollute the Cesàro average and should be kept aside.

2.3.3 Introduction of another chain, and stopping criterion

Our aim here is to build a reliable stopping criterion to S_n^b . The tuning of parameter b will be treated in the next subsection. Now we wish to know for which n the average S_n^b has reached \hat{u}_{LSE} with a given precision, say ε . In other words we wish to

$$\text{stop when } \|S_n^b - \hat{u}_{LSE}\| \leq \varepsilon \text{ holds.}$$

To this end, let us consider another Markov chain (\tilde{U}_n) defined as (U_n) and independent from it. Define \tilde{S}_n^b naturally by

$$\tilde{S}_n^b = \frac{1}{n-b} \sum_{k=b+1}^n \tilde{U}_k.$$

Both (S_n^b) and (\tilde{S}_n^b) converge towards \hat{u}_{LSE} with the same convergence rate. The comparison between (\tilde{S}_n^b) and (S_n^b) can help us measure the convergence of our approximation, as it will be seen in this paragraph. The expected distance between the two sequences (S_n^b) and (\tilde{S}_n^b) satisfies

$$\begin{aligned} \mathbb{E}\|\tilde{S}_n^b - S_n^b\|^2 &= \mathbb{E}\|\tilde{S}_n^b - \mathbb{E}[\tilde{S}_n^b]\|^2 + \mathbb{E}\|S_n^b - \mathbb{E}[S_n^b]\|^2 \\ &= 2 \text{Span}(S_n^b), \end{aligned}$$

i.e. the distance between (S_n^b) and (\tilde{S}_n^b) is twice the span of the estimator. This is because the chains (U_n) and (\tilde{U}_n) are independent with the same distribution.

Let us assume that the burn-in b is large enough to have a negligible bias $\|\mathbb{E}S_n^b - \hat{u}_{LSE}\|$ compared to the span term (2.11). Then the error of the estimator S_n^b is controlled by the distance between (S_n^b) and (\tilde{S}_n^b)

$$\mathbb{E}\|S_n^b - \hat{u}_{LSE}\|^2 \approx \text{Span}(S_n^b) = \frac{1}{2}\mathbb{E}\|\tilde{S}_n^b - S_n^b\|^2.$$

Once again, the images are supposed to be large and stationary in space, such that $\mathbb{E}\|S_n^b - \hat{u}_{LSE}\|^2 \approx \|S_n^b - \hat{u}_{LSE}\|^2$ and $\mathbb{E}\|\tilde{S}_n^b - S_n^b\|^2 \approx \|\tilde{S}_n^b - S_n^b\|^2$. Then

$$\|S_n^b - \hat{u}_{LSE}\| \approx \frac{1}{\sqrt{2}}\|\tilde{S}_n^b - S_n^b\|. \quad (2.18)$$

Moreover, notice that the mean $\frac{S_n^b + \tilde{S}_n^b}{2}$ is a better estimator than S_n^b or \tilde{S}_n^b taken separately, because it satisfies

$$\begin{aligned} \left\| \frac{S_n^b + \tilde{S}_n^b}{2} - \hat{u}_{LSE} \right\|^2 &= \left\| \frac{S_n^b - \hat{u}_{LSE}}{2} + \frac{\tilde{S}_n^b - \hat{u}_{LSE}}{2} \right\|^2 \\ &= \left\| \frac{S_n^b - \hat{u}_{LSE}}{2} \right\|^2 + \left\| \frac{\tilde{S}_n^b - \hat{u}_{LSE}}{2} \right\|^2 \\ &= \frac{1}{2}\|S_n^b - \hat{u}_{LSE}\|^2, \end{aligned} \quad (2.19)$$

such that, combining with Equation (2.18), we get the main result of this subsection:

Theorem 2.5 *Under the assumptions mentioned above (Central Limit Theorem for the chain, negligible bias, large and stationary images), we have*

$$\left\| \frac{S_n^b + \tilde{S}_n^b}{2} - \hat{u}_{LSE} \right\| \approx \frac{1}{2} \|\tilde{S}_n^b - S_n^b\|. \quad (2.20)$$

That is, the mean $\frac{S_n^b + \tilde{S}_n^b}{2}$ is twice closer to \hat{u}_{LSE} than S_n^b or \tilde{S}_n^b alone. This allows us to speed up the convergence for free. The two equations (2.18) and (2.20) are satisfied in practice as illustrated on Figure 2.2, provided the bias is negligible (i.e. the burn-in parameter b is large enough).

Our stopping criterion will then be the following

stop when $\ \tilde{S}_n^b - S_n^b\ \leq 2\varepsilon$, and return $\frac{S_n^b + \tilde{S}_n^b}{2}$,
--

because $\frac{S_n^b + \tilde{S}_n^b}{2}$ is expected to lie at a distance ε at most from \hat{u}_{LSE} .

2.4 Burn-in selection

As evoked in previous section, the number of iterations b being kept out from the average S_n^b (Equation (2.10)) is a determining factor for the convergence of the algorithm. An efficient

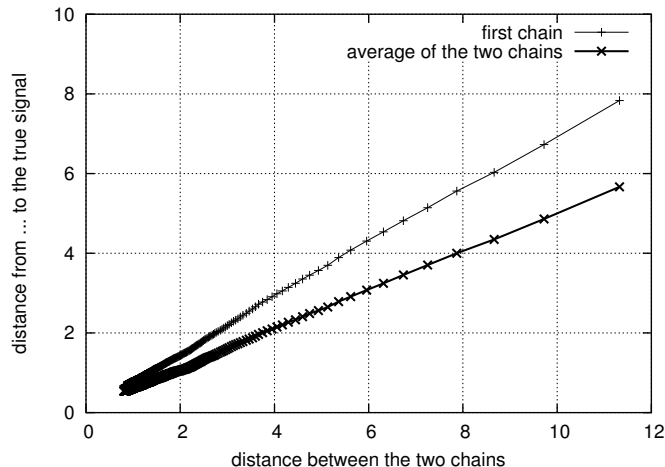


Figure 2.2: Convergence control. The target precision $\|S_n^b - \hat{u}_{LSE}\|$ is plotted against $\|\tilde{S}_n^b - S_n^b\|$ (thin line) for a fixed value of b , each little cross standing for a value of n . The relation is linear, which suggests that Eq. (2.18) holds quite precisely, and that the bias of the estimator S_n^b is negligible compared to its span, i.e. the chain has reached its stationary regime. The thick line represents $\|\frac{S_n^b + \tilde{S}_n^b}{2} - \hat{u}_{LSE}\|$ as a function of $\|\tilde{S}_n^b - S_n^b\|$. Again the relation is linear and is consistent with Eq. (2.20). Moreover the slope is smaller by a factor close to $1/\sqrt{2}$ than the one of the thin curve, as predicted by (2.19). This indicates that for a given number of iterations n , $\frac{S_n^b + \tilde{S}_n^b}{2}$ achieves a faster convergence towards \hat{u}_{LSE} than S_n^b alone. (Experiment made with $\lambda = 30$ and $\sigma = 10$ on a noisy image).

tuning of this burn-in parameter b can therefore be crucial. As the bias-variance decomposition (2.17) makes explicit, taking b too small make the average S_n^0 deviate from its target value \hat{u}_{LSE} for practical values of n (large bias, small variance). Conversely when b is very large then there may remain a too small number of terms in S_n^b which will spend more time to converge towards \hat{u}_{LSE} (small bias, large variance). This trade-off phenomenon is illustrated on Figure 2.3. We propose here a practical method for the tuning of b such that an optimum in the bias-variance decomposition is reached.

We look for a parameter b minimizing the error $\|\frac{S_n^b + \tilde{S}_n^b}{2} - \hat{u}_{LSE}\|$. However this quantity is not practically computable because \hat{u}_{LSE} is unknown. Hopefully as Equation (2.20) tells, it is equivalent to minimizing the distance between the average of the chains $\|S_n^b - \tilde{S}_n^b\|$, which is computable. Figure 2.3 demonstrates that the approximation

$$\arg \min_{b \in \{1, \dots, n\}} \left\| \frac{S_n^b + \tilde{S}_n^b}{2} - \hat{u}_{LSE} \right\| \approx \arg \min_{b \in \{1, \dots, n\}} \|S_n^b - \tilde{S}_n^b\|$$

holds in practice.

However for a given number of iterations n , finding the optimal burn-in b requires the knowledge of all the distances ($\|S_n^b - \tilde{S}_n^b\|$) for $1 \leq b \leq n$, and in practice the knowledge of all the images ($U_{|\Omega|b}$) and ($\tilde{U}_{|\Omega|b}$) for $1 \leq b \leq n$. For large values of n , it may not be possible to store all these images. That is why we design an efficient algorithm which only uses a limited number




Figure 2.3: Burn-in selection. The left graph illustrates the selection of the burn-in parameter b . For a fixed number of iterations n , the quantities $\|\frac{S_n^b + \tilde{S}_n^b}{2} - \hat{u}_{LSE}\|$ (thin curve) and $\|S_n^b - \tilde{S}_n^b\|$ (thick curve) are plotted against $b \in \{1, \dots, n\}$. The curves reflect the bias-variance trade-off: for small b , the convergence is not optimal because of a big bias, and for large b , the convergence is again suboptimal because of a large variance. The best burn-in given by $\arg \min_b \|\frac{S_n^b + \tilde{S}_n^b}{2} - \hat{u}_{LSE}\|$ can be well approximated by the \bar{b} that minimizes the other quantity $\|S_n^b - \tilde{S}_n^b\|$, because both quantities reach the optimum at the same value approximately. The right graph is about the proportionality of $\|\frac{S_n^b + \tilde{S}_n^b}{2} - \hat{u}_{LSE}\|$ and $\|S_n^b - \tilde{S}_n^b\|$. The ratio between the curves of the left graph is plotted. When the bias is negligible, this ratio should be 2 according to Eq. (2.20). We can see that for b larger than \bar{b} , the ratio is practically 2. This gives another clue to the fact that the bias is avoided when taking a burn-in equal to \bar{b} . (Experiment made with $\lambda = 30$ and $\sigma = 10$ on a noisy image).

of stored images, which approximately minimizes $\|S_n^b - \tilde{S}_n^b\|$.

Instead of taking

$$\bar{b} = \arg \min_{b \in \{1, \dots, n\}} \|S_n^b - \tilde{S}_n^b\|$$

where the minimization holds on all $\{1, \dots, n\}$, and returning

$$\frac{S_n^{\bar{b}} + \tilde{S}_n^{\bar{b}}}{2}$$

which would require the knowledge of all the images $(S_n^b)_{b \in \{1, \dots, n\}}$ and $(\tilde{S}_n^b)_{b \in \{1, \dots, n\}}$, we restrict b to live inside a geometric grid of \mathbb{N} of the form

$$\lfloor \lambda^{\mathbb{N}} \rfloor = \{\lfloor \lambda^k \rfloor, k \in \mathbb{N}\},$$

with $\lambda > 1$, and to be larger than a proportion p of n . Hence for a fixed n , we only store (S_n^b) and (\tilde{S}_n^b) with $b \in \lfloor \lambda^{\mathbb{N}} \rfloor \cap \{1, \dots, n\}$, and $b \geq pn$, which yields a maximum of $\frac{-2 \log p}{\log \lambda}$ images, which is independent of n .

We select b minimizing $\|S_n^b - \tilde{S}_n^b\|$ among the available data

$$\hat{b} = \arg \min_{b \in \lfloor \lambda^{\mathbb{N}} \rfloor \cap \{1, \dots, n\}, b \geq pn} \|S_n^b - \tilde{S}_n^b\|$$

and we return

$$\frac{S_n^{\hat{b}} + \tilde{S}_n^{\hat{b}}}{2}.$$

The distance between the optimal burn-in \bar{b} and its estimate \hat{b} can be bounded from above by

$$|\bar{b} - \hat{b}| \lesssim n \cdot \max\left(p, \frac{\lambda - 1}{\lambda}\right)$$

(p when $\bar{b} = 0$, and $(\lambda - 1)/\lambda$ when $\bar{b} = n - 1$ and $n \in \lfloor \lambda^{\mathbb{N}} \rfloor$). This bound measures the maximum number of additional iterations which should be computed when mistaking \bar{b} for \hat{b} . It is proportional to n , and the linear coefficient is close to 0 when $p \approx 0$ and $\lambda \approx 1$, which suggests that this method requires a small percentage at worst of extra iterations to be run.

In practice we take $\lambda = 1.2$ and $p = 1/6$. Then $\frac{-2 \log p}{\log \lambda} \approx 19.7$ and this yields at most 20 images (10 images S_n^b and 10 images \tilde{S}_n^b) to be stored, and a maximum of $n/6$ additional iterations to be computed (i.e. at most 17% extra iterations).

The interaction between the tuning of n and b is summed up in the Algorithm 3.

Algorithm 3 Burn-in algorithm

choose $\lambda = 1.2$, $p = 1/6$, $\varepsilon = 1$.
 start from initializations U_0 and \tilde{U}_0
 $S_0 = 0$, $\tilde{S}_0 = 0$
 $n = 0$
repeat
 $n \leftarrow n + 1$
 compute U_n from U_{n-1} as in Algorithm 2
 compute \tilde{U}_n from \tilde{U}_{n-1} as in Algorithm 2
 form $S_n = S_{n-1} + U_n$ and $\tilde{S}_n = \tilde{S}_{n-1} + \tilde{U}_n$
 if $n \in \lfloor \lambda^{\mathbb{N}} \rfloor$, $\left\{ \begin{array}{l} \text{store } S_n \text{ and } \tilde{S}_n \\ \text{erase the } S_k \text{ and } \tilde{S}_k \text{ with } k < pn \\ \text{set } \hat{b} = \arg \min_{b \in \lfloor \lambda^{\mathbb{N}} \rfloor} \left\| \frac{S_n - S_b}{n-b} - \frac{\tilde{S}_n - \tilde{S}_b}{n-b} \right\| \end{array} \right.$
until $\left\| \frac{S_n - S_{\hat{b}}}{n-\hat{b}} - \frac{\tilde{S}_n - \tilde{S}_{\hat{b}}}{n-\hat{b}} \right\| \leq 2\varepsilon$
return $\frac{(S_n - S_{\hat{b}}) + (\tilde{S}_n - \tilde{S}_{\hat{b}})}{2(n-\hat{b})}$.

2.5 Optimal scaling

Our model depends on the scaling α of the increment distribution q of (Y_n) , which we recall to be

$$q(u'|u) = \begin{cases} 0 & \text{if } u' \text{ and } u \text{ differ on 2 pixels or more,} \\ \frac{1}{2\alpha|\Omega|} \mathbb{1}_{[u(x)-\alpha, u(x)+\alpha]}(u'(x)) & \text{if } \exists x \in \Omega, \forall y \neq x, u'(y) = u(y). \end{cases}$$

The selection of this parameter α can be crucial, and a poor choice of α can lead to an extremely inefficient algorithm. Indeed for very small values of α , the increment distribution q proposes small jumps, most of which will be accepted, because the continuity of π implies that $\pi(u')/\pi(u) \approx 1$ when u and u' are nearby. But the size of the jumps is too small for the space to be rapidly explored. This is most true in higher dimensions. The Markov chain will then be slow to converge to the stationary distribution. Conversely for very large values of α the increment distribution q proposes very large jumps onto images u of very low probability $\pi(u)$ (because π is log-concave). These jumps will practically always be rejected, and the chain shall not move during a great number of iterations, slowing down the algorithm. This situation discloses the need for a correct tuning of α , and is illustrated on Figure 2.4.




Figure 2.4: Illustration for optimal scaling in a 2-D framework. (a) Target distribution $\pi \sim \mathcal{N}(0, \Gamma)$ with $\Gamma = (2, 1; 1, 1)$ ($Sp(\Gamma) \approx \{2.6, 0.4\}$). We consider the case where the increment distribution is an increment $\Upsilon_n + \alpha \mathcal{U}_{[-1,1]}$ on the random coordinate $i \in \{1, 2\}$. (b) Simulation of a MCMC with 10000 points with $\alpha = 0.02$. The chain is very slow to reach the stationary regime, and $\tau_{accept} = 99.6\%$. (c) Simulation of a MCMC with 10000 points with $\alpha = 200$. Only 68 iterations are accepted, and the chain does not manage to explore the stationary distribution; $\tau_{accept} = 0.68\%$. (d) Simulation of a MCMC with 10000 points with $\alpha = 2$. The stationary distribution is correctly explored by the chain; $\tau_{accept} = 57.2\%$.

As explained by G.O. Roberts and J.S. Rosenthal in [Roberts and Rosenthal 2001], the optimization of the scaling parameter α in random walk-based Hastings Metropolis algorithms can be considered in two possible ways. The first point of view is the minimization of the variance γ implied in the Central Limit Theorem, leading to an efficiency criterion

$$e_g = \left(\lim_{n \rightarrow \infty} \text{var} \left(\frac{1}{n} \sum_{k=1}^n g(U_k) \right) \right)^{-1}$$

which has to be maximized for all component functions $g = g_x$. However this measure of efficiency highly depends on the function g chosen: for two different Markov chains, different functions g can order their efficiency differently. As [Roberts and Rosenthal 2001] demonstrates, the high-dimensional limit case make the efficiency coefficients e_g less dependent from function g , and merge into a unique quantity independent from g , which is related to the acceptance rate

$$\tau_{accept} = \lim_{n \rightarrow \infty} \frac{1}{n} \#\{\text{accepted moves}\}$$

measuring the proportion of iterations where the proposal $U_{k+1/2}$ has been accepted (Algorithm 1). This is the second point of view which we will focus on here. Several theoretical results [Bédard 2007; Breyer and Roberts 2000; Neal and Roberts 2006; Roberts and Rosenthal 2001] developed in several frameworks in high dimension tend to prove that an optimal scaling α is reached when the acceptance rate reaches a universal constant

$$\tau_{\text{accept}} \approx 0.234.$$

The result which comes closest to our framework is probably given by [Breyer and Roberts 2000], where the stationary distribution π is assumed to be a Gibbs field defined on \mathbb{Z}^2 not subject to phase transition phenomenon. We transcribe it here informally; we do not aim at completeness and omit technical assumptions.

Theorem 2.6 [Breyer and Roberts 2000] *Take (A_i) a sequence of rectangular grids increasing to \mathbb{Z}^2 as $i \rightarrow \infty$. Consider the random walk-based Hastings-Metropolis algorithm run where the increment distribution is Gaussian with variance $\alpha_i^2 = l^2/|A_i|$. Consider the subsampled chain*

$$Z_t(i) = \Upsilon_{\lfloor t|A_i| \rfloor}(i),$$

i.e. the i -th component of the chain speeded up by a factor $|A_i|$. Then under certain conditions, the most important one being that the field's correlations must decrease exponentially fast as a function of distance (which insures the absence of phase transition in π), $Z(i)$ converges in distribution to a limiting infinite-dimensional diffusion process on $\mathbb{R}^{\mathbb{Z}^2}$. Moreover, the speed of this limiting diffusion can be written as a function of τ_{accept} and is proportional to

$$\tau_{\text{accept}} \left[\phi^{-1} \left(\frac{\tau_{\text{accept}}}{2} \right) \right]^2$$

and thus is maximal when $\tau_{\text{accept}} = 0.234$ (up to 3 decimal places).

[Neal and Roberts 2006] consider more general increment distributions such as partially updating distributions, but i.i.d. or Gaussian target distributions. [Bédard 2007] as well as [Breyer and Roberts 2000] considers non-i.i.d. target distributions but restricts the increment distribution to be Gaussian and i.i.d. Until now we found no result perfectly adapting to our case.

Theorem 2.6 means that in high dimension α should be tuned in order to reach the acceptance rate 0.234. This optimal acceptance rate seems low: more than 3 proposals out of 4 are rejected. But a smaller acceptance rate means a larger scale and a faster exploration of the support of π . Besides, a scaling factor α should be chosen proportional to $1/\sqrt{|\Omega|}$ where $|\Omega|$ is the number of pixels in the considered images, i.e. the problem's dimension.

Remark 1: An insight study of the convergence rate in function of τ_{accept} shows that τ_{accept} in $[0.1, 0.5]$ suffices to have a correct convergence rate (the efficiency curve is smooth at its maximum $\tau_{\text{accept}} = 0.234$) [Roberts and Rosenthal 2001].

Remark 2: In low dimension spaces [Gelman et al. 1996], better rates of convergence are achieved for $\tau_{\text{accept}} \approx 0.5$. Hence in the illustrating example given in Figure 2.4 the scaling value $\alpha = 2$ is almost optimal.

In our experiments the scale α is initialized by the gradient average of the datum image v

$$\alpha_0 = \frac{1}{|\Omega|} \sum_{x \in \Omega} |\nabla v|(x) \quad (2.21)$$

(typical deviation between neighbor pixels) and is adjusted throughout the iterations by a simple dichotomous procedure summed up in Algorithm 4.

Algorithm 4 Optimal scaling algorithm

```

set  $\alpha = \alpha_0$ ,  $\alpha_{min} = 0$ ,  $\alpha_{max} = \max_x v(x) - \min_x v(x)$ 
 $n = 0$ 
repeat
  run the chain from  $\Upsilon_{|\Omega|n+1}$  to  $\Upsilon_{|\Omega|(n+1)}$ 
  compute  $\tau_{accept} = \frac{1}{|\Omega|} \#\{\text{accepted moves between iterations } |\Omega|n+1 \text{ and } |\Omega|(n+1)\}$ 
until  $\tau_{accept}$  is stabilized.
repeat
  run the chain from  $\Upsilon_{|\Omega|n+1}$  to  $\Upsilon_{|\Omega|(n+1)}$ 
  compute  $\tau_{accept} = \frac{1}{|\Omega|} \#\{\text{accepted moves between iterations } |\Omega|n+1 \text{ and } |\Omega|(n+1)\}$ 
  if  $\tau_{accept} < 0.23$  set  $\begin{cases} \alpha_{max} \leftarrow \alpha \\ \alpha \leftarrow \frac{\alpha + \alpha_{min}}{2} \end{cases}$ 
  if  $\tau_{accept} > 0.25$  set  $\begin{cases} \alpha_{min} \leftarrow \alpha \\ \alpha \leftarrow \frac{\alpha + \alpha_{max}}{2} \end{cases}$ 
   $n \leftarrow n + 1$ 
until  $\tau_{accept} \in [0.23, 0.25]$ 
return  $\alpha$ .
```

Remark 3: The chain must be in its stationary regime, i.e. Algorithm 4 should be run after a preliminary burn-in phase. Moreover the algorithm should be run before the first iterations are taken inside the sum S_n^b . Indeed if α is still varying during the burn-in period, then the Markov chain is heterogeneous (the transition distribution depends on the iteration number), and the convergence theorems cannot be applied. However some recent papers such as [Andrieu and Moulines 2006] prove Ergodic and Central Limit Theorems for some cases of scaling adaptive Markov chains.

Remark 4: The optimal scaling factor α obviously depends on the parameters λ and σ of π (2.4). When σ is large (i.e. π has a large bandwidth), the chain needs to be more mobile to explore a larger support of π , and α is then large. Similarly,

$$\begin{cases} \sigma \nearrow \Rightarrow \alpha \nearrow & \text{when } \lambda \text{ is fixed,} \\ \lambda \nearrow \Rightarrow \alpha \searrow & \text{when } \sigma \text{ is fixed,} \\ \sigma \nearrow \Rightarrow \alpha \nearrow & \text{when } \frac{\lambda}{\sigma^2} \text{ is fixed.} \end{cases}$$

Hence the Algorithm 4 is usually necessary to refine the initialization α_0 (2.21) and to take λ and σ into consideration.




Figure 2.5: Convergence speed as a function of the scale parameter α of q . For α too small or too large the quantity $\log \|\tilde{S}_n^b - S_n^b\|$ decreases to $-\infty$ quite slowly, while for middle values of α the convergence rate is optimal (here for $\alpha = 80$). The proposed method returns $\alpha \approx 40$, which is quite far away from the value 80; but the convergence rate is almost the same, which matters here. (Experiments made with a 64×64 noisy image and $\lambda = 20$ and $\sigma = 10$).

Algorithm 4 has been tested, and the output $\hat{\alpha}$ (such that $\tau_{accept} \approx 0.234$, and supposed to maximize the convergence speed of the MCMC algorithm) is compared with other scaling factors. As it can be seen on Figure 2.5, the algorithm returns a value of α which has the same order of magnitude as the one optimizing the speed of convergence ($\hat{\alpha} = 40$ instead of $\alpha_{opt} = 80$). Its relative lack of precision is balanced by the fact that the rate of convergence is not very sensitive to α at its maximum. In other words, any $\alpha \in (40, 160)$ achieves a correct convergence rate, and the estimate $\hat{\alpha}$ is enough for our task.

2.6 Numerical variants – why we won’t run several chains

2.6.1 Several independent MCMCs

A few pages before, we introduced a second MCMC to control the convergence of the algorithm. We noticed that averaging the two MCMCs gave a better precision, because when neglecting the bias term we got

$$\left\| \frac{S_n^b + \tilde{S}_n^b}{2} - \hat{u}_{LSE} \right\| \approx \frac{1}{\sqrt{2}} \|S_n^b - \hat{u}_{LSE}\|.$$

We could believe that using other chains would again improve the result. If we run N chains $(U^{(1)}, \dots, U^{(N)})$, then still omitting the bias term, the average of all the chains is much better

than one chain alone, i.e. with obvious notations

$$\left\| \frac{1}{N} \sum_{k=1}^N S_n^{b,k} - \hat{u}_{LSE} \right\| \approx \frac{1}{\sqrt{N}} \|S_n^b - \hat{u}_{LSE}\|.$$

Besides the proximity of all the chains $(S_n^{b,k})_{1 \leq k \leq N}$ could be used as a more precise convergence criterion than $\|S_n^b - \tilde{S}_n^b\|$ only. But it should be kept in mind that running N chains costs N times the price of one chain. In other words, for a given time complexity, the N chains should be N times shorter. As empirically commented by [Geyer 1992], running several chains which are too short cannot be used as a convergence criterion; conversely, when the chains are long enough, one run suffices. Several short chains are then less efficient than a unique long chain, as will illustrate Propositions 2.6 and 2.7.

The quadratic error related the estimation of \hat{u}_{LSE} by N chains of length n with uniform burn-in b can be computed and bounded by

$$\mathbb{E} \left\| \frac{1}{N} \sum_{k=1}^N S_n^{b,k} - \hat{u}_{LSE} \right\|^2 = \mathbb{E} \left\| \frac{1}{N} \sum_{k=1}^N S_n^{b,k} - \mathbb{E} \frac{1}{N} \sum_{k=1}^N S_n^{b,k} \right\|^2 + \left\| \mathbb{E} \left[\frac{1}{N} \sum_{k=1}^N S_n^{b,k} \right] - \hat{u}_{LSE} \right\|^2 \quad (2.22)$$

$$= \text{Span} \left(\frac{1}{N} \sum_{k=1}^N S_n^{b,k} \right) + \left\| \mathbb{E} S_n^{b,1} - \hat{u}_{LSE} \right\|^2$$

$$= \frac{1}{N} \text{Span}(S_n^{b,1}) + \left\| \mathbb{E} S_n^{b,1} - \hat{u}_{LSE} \right\|^2 \quad (2.23)$$

$$\leq \frac{A}{N(n-b)} + \frac{B\delta^{2|\Omega|b}}{(n-b)^2}. \quad (2.24)$$

Indeed, (2.22) is the bias-variance decomposition of the estimator $\frac{1}{N} \sum_{k=1}^N S_n^{b,k}$, (2.23) holds because the $(S_n^{b,k})_k$ are i.i.d., and (2.24) is a consequence of the central limit theorem and of the geometric ergodicity of the chains (see Equation (2.17)).

Next proposition shows that for a given burn-in b , a unique long chain with size Nn performs better than N chains with size n . Let

$$e(n, N, b) = \frac{A}{N(n-b)} + \frac{B\delta^{2|\Omega|b}}{(n-b)^2}$$

denote the upper bound of the quadratic error related to the running of N chains with length n with burn-in b (uniform over the N chains).

Proposition 2.6 *For any $A, B > 0$, any $\delta \in (0, 1)$, any $N, n \in \mathbb{N}^*$, and any $b \in \{0, \dots, n-1\}$ we get*

$$e(Nn, 1, b) \leq e(n, N, b).$$

In other words, the error bound is tighter for one chain with size Nn than for N chains with size n , when the burn-in parameter is fixed.

Proof : As

$$\frac{A}{Nn - b} \leq \frac{A}{N(n - b)} \quad \text{and} \quad \frac{B\delta^{2|\Omega|b}}{(Nn - b)^2} \leq \frac{B\delta^{2|\Omega|b}}{(n - b)^2},$$

the inequality holds. \square

This result can be refined to the case where the burn-in minimizes the error $e(n, N, b)$. Indeed, one expects the optimal burn-in of a long chain to be larger than for many short chains. Next result states that the error bound minimized with respect to b is still smaller for the long chain than for N short chains.

Proposition 2.7 *For any $A, B > 0$, any $\delta \in (0, 1)$ and any $N, n \in \mathbb{N}^*$, we get*

$$\min_{0 \leq b < Nn} e(Nn, 1, b) \leq \min_{0 \leq b < n} e(n, N, b).$$

Proof : Fix $\delta \in (0, 1)$, N and n in \mathbb{N}^* . Proposition 2.6 proves that for all $b \in \{0, \dots, n - 1\}$,

$$e(Nn, 1, b) \leq e(n, N, b).$$

This inequality can be minimized with respect to $b \in \{0, \dots, n - 1\}$ on both sides, which yields

$$\min_{0 \leq b < n} e(Nn, 1, b) \leq \min_{0 \leq b < n} e(n, N, b).$$

Now, as

$$\min_{0 \leq b < Nn} e(Nn, 1, b) \leq \min_{0 \leq b < n} e(Nn, 1, b),$$

we get the desired result. \square

Both bias and variance are worse when considering more than one chain, even if the burn-in is optimized: we lose on both counts.

2.6.2 Interacting MCMCs – Sampling Importance Resampling

We consider a more complex procedure than independent Hastings-Metropolis chains, known as *Sampling Importance Resampling* (SIR) [Gordon et al. 1993] (see [Arulampalam et al. 2002] for an overview). Instead of running N chains independently, we make them interact from times to times. The resampling step allows the low-energy chains to be duplicated and high-energy chains to be killed, in order to conquer the support of π faster.

More precisely, one iteration of the SIR algorithm is as follows:

- run one iteration on all the chains $(U^{(1)}, \dots, U^{(N)})$.
- compute the vector of pointwise densities $(\pi(U^{(1)}), \dots, \pi(U^{(N)}))$ and normalize it, to obtain a vector of normalized weights, here called *empirical weights*

$$\pi_{emp}(U^{(k)}) = \frac{1}{\sum_{l=1}^N \pi(U^{(l)})} \pi(U^{(k)}).$$

(There is no need to know the normalizing factor in π). The weight vector π_{emp} is a discrete approximation of the posterior distribution.

- compute the effective number of chains

$$N_{eff} = \frac{1}{\sum_{k=1}^N (\pi_{emp}(U^{(k)}))^2}.$$

- if N_{eff} is less than a given threshold N_{thr} , then resample the chains $(U^{(1)}, \dots, U^{(N)})$ with respect to the discrete distribution given by the weights $\pi_{emp}(U^{(k)})$ (this step is known as *resampling*). Chains $U^{(k)}$ with large weight are expected to be duplicated while the chains with small weight are likely to be killed.
- At the end of the procedure return the global average of the iterations on all the chains

$$\frac{1}{N(n-b)} \sum_{k=1}^N \sum_{i=b+1}^n U_i^{(k)}.$$

In practice the resampling step is crucial because it circumvents a current degeneracy problem, where the empirical weight vector often converges to a Dirac. When one chain has a much higher posterior probability than all the others, then the effective number of chains is close to zero and resampling is desirable.

However in our TV-LSE framework, we believe that resampling would bring poor (if any) improvement on the convergence speed of the algorithm, due to the high dimensionality of the problem and to the unimodality of the target distribution. Numerical experiments show that the degeneracy problem does not occur in our case, because the empirical weight vector quickly converges to a uniform weight vector. If resampling yet occurred (if N_{thre} is very close to N for instance), it would be very close to a uniform resampling, and the SIR algorithm would then be almost equivalent to running several independent MCMCs.

Indeed, as shows Figure 2.6 (left), a MCMC started with a uniform white random noise on $[0, 255]$ has an energy which practically does not depend on its initialization and on the realization of the chain. In other words, for a given iteration n , the energy $E_\lambda(U_n^{(k)})$ is almost independent of k , and the discrete importance distribution π_{emp} is not far from being uniform. In order to quantify the closeness of π_{emp} to a uniform weight $\pi_{unif}(U^{(k)}) = \frac{1}{N}$, we used the entropy of π_{emp} , given by

$$H(\pi_{emp}) = - \sum_{k=1}^N \log_2(\pi(U^{(k)}))\pi(U^{(k)}).$$

The entropy is maximal and equals $\log_2(N)$ if and only if $\pi_{emp} = \pi_{unif}$. We plotted $H(\pi_{emp})$ against the number of iterations on Figure 2.6 (right). When n increases, the curve rapidly converges towards $\log_2(N)$. This indicates that the importance resampling step is essentially a uniform sampling on the N chains, certainly because of the high dimensionality. The procedure is thus almost equivalent to running N independent MCMCs and summing them together. But this is suboptimal, because of Proposition 2.7). That is why in our experiments (Chapter 4) we run the minimal number of chains, i.e. 2 chains, to be able to apply our stopping criterion (see Section 2.3).

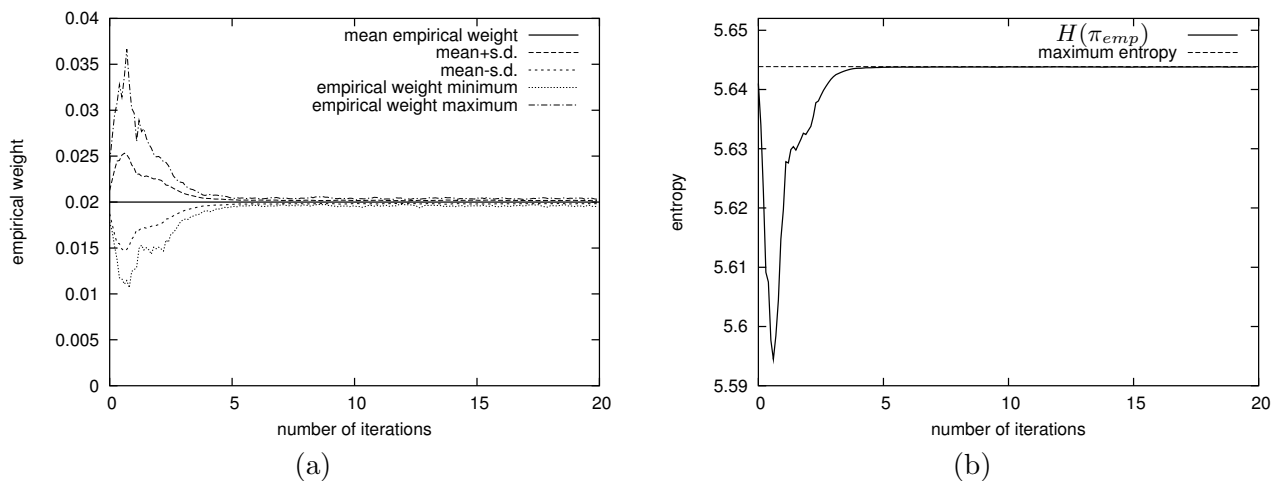


Figure 2.6: Resampling in SIR on several chains is almost uniform resampling. (a) Some features of π_{emp} computed on a i.i.d. sample of 50 MCMCs are plotted against the number n of iteration (1 iteration means $|\Omega|$ runs in the Hastings-Metropolis algorithm). First the average weight $\pi_{emp}(U^{(k)})$ of the 50 chains ($U^{(k)}$) is of course constant and equal to $1/50$. The standard deviation has also been computed against n , and occurs to be very small compared to the average as soon as a small number of iterations have been run. Indeed, the sum of the average line with the s.d. (2nd curve) and the difference of the two (3rd curve) happen to be very close to the average line as soon as 4 or 5 iterations are reached. Moreover the minimum and the maximum of $\pi_{emp}(U^{(k)})$ converge to $1/50$ within very few (5 or 6) iterations. This graph demonstrates that π_{emp} is not far from being uniform. The values have been computed from a sample of 50 i.i.d. chains with size 64×64 started with a uniform white random noise on $[0, 255]$. (b) To make the closeness of π_{emp} to a uniform weight more accurate, the entropy of the empirical weight distribution is plotted against the number n of iteration. This entropy rapidly converges to $\log_2(50) \approx 5.64$, which corresponds to the maximum entropy attained by a uniform weight distribution. Note that the y -scale which is very concentrated around $\log_2(50)$. This means that after 4 iterations importance resampling is almost uniform and the interaction between chains is non-existent. The empirical distribution has been computed on the 50 i.i.d. chains of (a).

Chapter 3

Theoretical properties

Résumé Dans ce chapitre, nous présentons quelques résultats théoriques concernant l'opérateur de débruitage TV-LSE. Tout d'abord, de petites propriétés d'invariance et de différentiabilité sont données. Puis nous donnons des résultats asymptotiques sur le comportement d'une image lorsque les paramètres du modèle sont poussés vers des valeurs limites. Ensuite, un théorème sur l'absence de *staircasing* dans le débruitage TV-LSE est démontré; il justifie a posteriori l'approche menée dans le chapitre 1. Ces résultats sont vrais pour de nombreux schémas numériques consistants avec la variation totale. Enfin, nous étudions le débruitage LSE d'un point de vue analytique, dans un cadre un peu plus général que TV. Nous montrons qu'il vérifie quelques propriétés de stabilité, et que sous certaines conditions, il est exactement équivalent à un débruitage MAP, à condition de considérer une loi a priori plus régulière pour le MAP. Cela laisse à penser que le résultat d'absence de *staircasing* pourrait se généraliser à d'autres modèles que le cadre TV.

Abstract In this chapter we present several theoretical results concerning the TV-LSE denoising operator. First, basic properties are given, among which invariance and differentiability results. Then we give asymptotic results on the denoised image when the parameters go to extreme values. Afterwards a theorem stating the absence of staircasing in TV-LSE denoised images is proven. This theorem justifies the construction of TV-LSE *a posteriori* and answers to our expectations raised in Chapter 1. All these results hold under quite loose conditions on the numerical scheme for the total variation. Last, we set up an analytic framework in which LSE denoising is proven to satisfy several stability results. LSE and MAP denoisings are then compared. Under some conditions upon the prior which are satisfied by the TV prior, it is established that a LSE denoiser is equivalent to a MAP denoiser built on a smoother prior. This gives another clue for the absence of staircasing.

Notations and assumptions We consider images defined on a finite set $\Omega \subset \mathbb{Z}^2$. The cardinality of Ω is denoted $|\Omega|$. For convenience let S_{LSE} denote the LSE denoising operator

and S_{MAP} the MAP denoising operator, i.e.

$$\begin{aligned} S_{LSE} : \mathbb{R}^\Omega &\longrightarrow \mathbb{R}^\Omega \\ v &\mapsto \frac{\int_{\mathbb{R}^\Omega} u e^{-\frac{\|u-v\|^2 + \lambda TV(u)}{2\sigma^2}} du}{\int_{\mathbb{R}^\Omega} e^{-\frac{\|u-v\|^2 + \lambda TV(u)}{2\sigma^2}} du} \end{aligned} \quad (3.1)$$

and

$$\begin{aligned} S_{MAP} : \mathbb{R}^\Omega &\longrightarrow \mathbb{R}^\Omega \\ v &\mapsto \arg \min_{u \in \mathbb{R}^\Omega} \|u - v\|^2 + \lambda TV(u). \end{aligned} \quad (3.2)$$

Therefore $S_{LSE}(v) = \hat{u}_{LSE}$, and $S_{MAP}(v) = \hat{u}_{MAP}$.

The gray level of an image u at pixel x will be equivalently denoted $u(x)$ or u_x , as convenient. We assume the space of images \mathbb{R}^Ω to be equipped with the inner product

$$\langle u, v \rangle = \frac{1}{|\Omega|} \sum_{x \in \Omega} u(x)v(x).$$

The norm $\|\cdot\|$ over \mathbb{R}^Ω will always denote the norm derived from this inner product, i.e. the classical Euclidean norm divided by the size of the image domain.

For a given v , E_λ will denote the Rudin-Osher-Fatemi energy

$$E_\lambda(u) = \|u - v\|^2 + \lambda TV(u).$$

The TV operator need not be specified in all the results of this chapter. That is why we simply assume that it is a non-negative mapping from $\mathbb{R}^\Omega \rightarrow \mathbb{R}^+$. Extra assumptions will be made explicit when needed.

3.1 Basic properties

Let us state several basic properties of the TV-LSE denoiser that could predict some behavior of the denoiser. We start by giving properties that LSE and MAP denoising have in common, and keep going on properties that are specific to LSE denoising.

3.1.1 Average preservation

In this subsection we assume TV shift-invariant, i.e. such that

$$\forall c \in \mathbb{R}, \forall u \in \mathbb{R}^\Omega, \quad TV(u + c) = TV(u).$$

Proposition 3.1 *For any image u let $\bar{u} = \frac{1}{|\Omega|} \sum_{x \in \Omega} u(x)$ denote the average gray level of u . Then for every $v \in \mathbb{R}^\Omega$,*

$$\overline{S_{LSE}(v)} = \bar{v}.$$

Proof : Let

$$\mathcal{E}_0 = \{u \in \mathbb{R}^\Omega : \bar{u} = 0\}$$

denote the subspace of images with mean zero. Splitting up the variable $u = u_0 + t\mathbf{1}$ into a zero-mean image $u_0 \in \mathcal{E}_0$ and a shift t (note that $\bar{u} = t = \langle u, \mathbf{1} \rangle$), the LSE-denoised image writes

$$\hat{u}_{LSE} = \frac{1}{Z} \int_{u_0 \in \mathcal{E}_0} \int_{t \in \mathbb{R}} u_0 e^{-\frac{E_\lambda(u_0+t\mathbf{1})}{2\sigma^2}} dt du_0 + \frac{1}{Z} \int_{u_0 \in \mathcal{E}_0} \int_{t \in \mathbb{R}} t e^{-\frac{E_\lambda(u_0+t\mathbf{1})}{2\sigma^2}} dt du_0 \mathbf{1}. \quad (3.3)$$

The first integral in Equation (3.3) has mean zero since it is a weighted average of zero-mean images. Let us focus on the second integral. The energy may be written

$$E_\lambda(u_0 + t\mathbf{1}) = \|\mathbf{1}\|^2 \left(t - \frac{\langle v - u_0, \mathbf{1} \rangle}{\|\mathbf{1}\|^2} \right)^2 + \|u_0 - v\|^2 - \frac{\langle \mathbf{1}, u_0 - v \rangle^2}{\|\mathbf{1}\|^2} + \lambda TV(u),$$

with $\|\mathbf{1}\| = 1$. Then integrating along t yields

$$\int_{\mathbb{R}} t e^{-\frac{(t - \langle v - u_0, \mathbf{1} \rangle)^2}{2\sigma^2}} dt = \sqrt{2\pi}\sigma \langle v - u_0, \mathbf{1} \rangle.$$

Hence the second integral in Equation (3.3) derives into

$$\begin{aligned} \frac{1}{Z} \int_{u_0 \in \mathcal{E}_0} \int_{t \in \mathbb{R}} t e^{-\frac{E_\lambda(u_0+t\mathbf{1})}{2\sigma^2}} dt du_0 &= \frac{1}{Z'} \int_{u_0 \in \mathcal{E}_0} \langle v - u_0, \mathbf{1} \rangle e^{-\frac{\|u_0 - v\|^2 - \langle \mathbf{1}, u_0 - v \rangle^2 + \lambda TV(u)}{2\sigma^2}} du_0 \\ &= \langle v, \mathbf{1} \rangle - \left\langle \frac{1}{Z'} \int_{u_0 \in \mathcal{E}_0} u_0 e^{-\frac{\|u_0 - v\|^2 - \langle \mathbf{1}, u_0 - v \rangle^2 + \lambda TV(u)}{2\sigma^2}} du_0, \mathbf{1} \right\rangle \\ &= \bar{v}, \end{aligned}$$

for the integral inside the inner product is again a weighted average of zero-mean images. Finally \hat{u}_{LSE} is a zero-mean image shifted by \bar{v} , therefore \hat{u}_{LSE} and v have equal average. \square

MAP denoising is also average preserving, because the addition of a constant to an image does not change the total variation, while $C \mapsto \|(v + C) - u\|^2$ is minimum when $\bar{v} + C = \bar{u}$ for any images u and v [Aubert and Kornprobst 2006].

3.1.2 Invariances

In this subsection several invariance properties of S_{LSE} are given. These properties are shared with MAP denoising.

Proposition 3.2 *Let $s : \mathbb{R}^\Omega \rightarrow \mathbb{R}^\Omega$ be a linear isometry such that for all $u \in \mathbb{R}^\Omega$, $TV \circ s(u) = TV(u)$ holds. Then*

$$\forall v \in \mathbb{R}^\Omega, \quad S_{LSE} \circ s(v) = s \circ S_{LSE}(v).$$

Proof : The change of variable $u' = s^{-1}(u)$ in the numerator and the denominator of (3.1) yields

$$S_{LSE}(s(v)) = \frac{\int s(u') e^{\frac{\|s(u') - s(v)\|^2 + \lambda TV(s(u'))}{2\sigma^2}} du'}{\int e^{\frac{\|s(u') - s(v)\|^2 + \lambda TV(s(u'))}{2\sigma^2}} du'},$$

because s being an isometry implies $ds(u') = du'$. Furthermore s is isometric, so we have $\|s(u') - s(v)\|^2 = \|u' - v\|^2$, and $TV(s(u')) = TV(u')$ thus

$$S_{LSE} \circ s(v) = \frac{\int s(u') e^{\frac{\|u'-v\|^2 + \lambda TV(u')}{2\sigma^2}} du'}{\int e^{\frac{\|u'-v\|^2 + \lambda TV(u')}{2\sigma^2}} du'} = s(S_{LSE}(v)).$$

because s is linear. □

Corollary 3.1 1. (**shift invariance**) *If TV satisfies*

$$\forall c \in \mathbb{R}, \forall u \in \mathbb{R}^\Omega, \quad TV(u + c) = TV(u),$$

then

$$\forall u \in \mathbb{R}^\Omega, \forall c \in \mathbb{R}, \quad S_{LSE}(u + c) = S_{LSE}(u) + c.$$

2. (**translation invariance**) *Assume here that Ω is a torus. Let $t \in \mathbb{Z}^2$. If τ_t is a translation operator defined by $\tau_t \circ u(x) = u(x - t)$ for all $u \in \mathbb{R}^\Omega$, and if TV satisfies*

$$TV \circ \tau_t = TV,$$

then

$$S_{LSE} \circ \tau_t = \tau_t \circ S_{LSE}.$$

3. (**$\pi/2$ -rotation invariance**) *If ρ is a $\pi/2$ -rotation sending Ω onto itself, and if TV satisfies*

$$TV \circ \rho = TV(u),$$

then

$$S_{LSE} \circ \rho = \rho \circ S_{LSE}.$$

4. (**“involution invariance”**) *If $s : \mathbb{R}^\Omega \rightarrow \mathbb{R}^\Omega$, a linear isometry, and $v \in \mathbb{R}^n$ are such that*

$$\begin{cases} s^2 = Id, \\ \forall u \in \mathbb{R}^\Omega, \quad TV(s(u)) = TV(u), \\ s(v) = v \end{cases}$$

then $s(\hat{u}_{LSE}) = \hat{u}_{LSE}$.

Proof : The operators s involved in the four examples are all linear and isometric, and satisfy $TV(s(u)) = TV(u)$ for all $u \in \mathbb{R}^\Omega$. Then Proposition 3.2 applies. □

These properties can help find the structure of the LSE denoised version of images containing lots of redundancies and structure. For example if TV is translation invariant, and if v is a constant image, then $S_{LSE}(v) = v$. Indeed, v is invariant under the translations of vector $(1, 0)$ and $(0, 1)$, and so is $S_{LSE}(v)$; moreover the average gray level of $S_{LSE}(v)$ is the same as v .

Finally $S_{LSE}(v)$ is a constant equal to v . Another example is given by the checkerboard, defined by

$$\exists x, y \in \mathbb{R}, \quad v(i, j) = \begin{cases} x & \text{if } i + j \text{ even} \\ y & \text{if } i + j \text{ odd.} \end{cases}$$

It is quite easy to see that $S_{LSE}(v)$ is also a checkerboard (use the invariance by translations of vectors $(1, 1)$ and $(1, -1)$) (but it seems difficult to get the parameters (x, y) of the denoised image). It means that S_{LSE} is likely to preserve the strong geometric features of images.

3.1.3 Regularity

Until now the properties of the LSE denoiser were common with the MAP denoiser. We now come to the first result that distinguishes LSE from MAP, and concerns the regularity of S_{LSE} .

Proposition 3.3 *The TV-LSE denoiser S_{LSE} is continuously differentiable.*

Note that the MAP denoiser is continuous and even Lipschitz continuous, but not differentiable. This can be seen by considering the family of input images

$$v^x(i, j) = \begin{cases} x & \text{if } i + j \text{ even} \\ -x & \text{if } i + j \text{ odd} \end{cases}$$

(dilated centered checkerboards), for which the MAP denoised image is the checkerboard

$$S_{MAP}(v^x)(i, j) = \begin{cases} \max(x - \lambda, 0) & \text{if } i + j \text{ even} \\ -\max(x - \lambda, 0) & \text{if } i + j \text{ odd.} \end{cases}$$

The mapping $x \mapsto S_{MAP}(v^x)(i, j)$ is not differentiable everywhere. Hence S_{MAP} is not differentiable.

Proof of Proposition 3.3 : First, note that if h is a unit vector of \mathbb{R}^Ω then

$$|t| < \varepsilon \quad \Rightarrow \quad \|u - v - th\|^2 \geq \frac{1}{2}\|u - v\|^2 - \varepsilon^2. \quad (3.4)$$

Let us first prove the continuity of S_{LSE} . Let $v \in \mathbb{R}^\Omega$ and $\varepsilon > 0$, and let us prove that S_{LSE} is continuous on the ball $B(v, \varepsilon)$. Let $v' \in B(v, \varepsilon)$. Then for both mappings $g(u) = 1$ and $g(u) = u$, the functions $v' \mapsto g(u) \exp(-\frac{\|u-v'\|^2 + \lambda TV(u)}{2\sigma^2})$ are continuous, and have a uniform upper bound thanks to (3.4) given by

$$\left| g(u) e^{-\frac{\|u-v'\|^2 + \lambda TV(u)}{2\sigma^2}} \right| \leq |g(u)| e^{-\frac{\frac{1}{2}\|u-v\|^2 - \varepsilon^2 + \lambda TV(u)}{2\sigma^2}} \in L^1(\mathbb{R}^\Omega),$$

which is an upper bound independent of $v' \in B(v, \varepsilon)$. Hence the continuity theorem under the integral sign can be applied, and both the numerator and denominator of S_{LSE} are continuous on $B(v, \varepsilon)$. Since the denominator cannot vanish, we deduce that S_{LSE} is continuous on \mathbb{R}^Ω .

To prove the differentiability of S_{LSE} , let h be a unit vector of \mathbb{R}^Ω and $\varepsilon > 0$. The functions

$$t \in (-\varepsilon, \varepsilon) \mapsto g(u) e^{-\frac{\|u-v-th\|^2 + \lambda TV(u)}{2\sigma^2}},$$

where $g(u) = 1$ or u , are \mathcal{C}^1 , with derivative

$$t \in (-\varepsilon, \varepsilon) \mapsto g(u) \left(-\frac{2t - 2\langle u-v, h \rangle}{2\sigma^2} \right) e^{-\frac{\|u-v-th\|^2 + \lambda TV(u)}{2\sigma^2}},$$

and satisfy thanks to (3.4)

$$\left| g(u) \left(-\frac{2t - 2\langle u-v, h \rangle}{2\sigma^2} \right) e^{-\frac{\|u-v-th\|^2 + \lambda TV(u)}{2\sigma^2}} \right| \leq |g(u)| \frac{\varepsilon + \|u-v\|}{\sigma^2} e^{-\frac{\frac{1}{2}\|u-v\|^2 + \lambda TV(u)}{2\sigma^2}} e^{\frac{\varepsilon^2}{2\sigma^2}}.$$

This bound is independent of t (provided that $|t| < \varepsilon$) and $h \in B(0, 1)$, and is integrable with respect to $u \in \mathbb{R}^\Omega$ since the Gaussian distribution admits moments of order 1 and 2. Now thanks to the derivation theorem under the integral sign, the denominator and the numerator of $S_{LSE}(v+th)$ satisfy

$$\begin{cases} \left. \frac{\partial}{\partial t} \int e^{-\frac{\|u-v-th\|^2 + \lambda TV(u)}{2\sigma^2}} du \right|_{t=0} = \int \frac{\langle u-v, h \rangle}{\sigma^2} e^{-\frac{\|u-v\|^2 + \lambda TV(u)}{2\sigma^2}} du \\ \left. \frac{\partial}{\partial t} \int u e^{-\frac{\|u-v-th\|^2 + \lambda TV(u)}{2\sigma^2}} du \right|_{t=0} = \int u \frac{\langle u-v, h \rangle}{\sigma^2} e^{-\frac{\|u-v\|^2 + \lambda TV(u)}{2\sigma^2}} du. \end{cases}$$

Therefore the differential of S_{LSE} writes

$$\begin{aligned} dS_{LSE}(v)(h) &= \frac{\int u \langle u-v, h \rangle e^{-\frac{E_\lambda(u)}{2\sigma^2}} du}{\int e^{-\frac{E_\lambda(u)}{2\sigma^2}} du} - \frac{\int u e^{-\frac{E_\lambda(u)}{2\sigma^2}} du}{\int e^{-\frac{E_\lambda(u)}{2\sigma^2}} du} \frac{\int \langle u-v, h \rangle e^{-\frac{E_\lambda(u)}{2\sigma^2}} du}{\int e^{-\frac{E_\lambda(u)}{2\sigma^2}} du} \\ &= \frac{\int u \langle u, h \rangle e^{-\frac{E_\lambda(u)}{2\sigma^2}} du}{\int e^{-\frac{E_\lambda(u)}{2\sigma^2}} du} - \frac{\int u e^{-\frac{E_\lambda(u)}{2\sigma^2}} du}{\int e^{-\frac{E_\lambda(u)}{2\sigma^2}} du} \frac{\int \langle u, h \rangle e^{-\frac{E_\lambda(u)}{2\sigma^2}} du}{\int e^{-\frac{E_\lambda(u)}{2\sigma^2}} du}. \end{aligned}$$

Thus

$$(dS_{LSE}(v))_{i,j} = \frac{\int u_i u_j e^{-\frac{\|u-v\|^2 + \lambda TV(u)}{2\sigma^2}} du}{\int e^{-\frac{\|u-v\|^2 + \lambda TV(u)}{2\sigma^2}} du} - \frac{\int u_i e^{-\frac{\|u-v\|^2 + \lambda TV(u)}{2\sigma^2}} du}{\int e^{-\frac{\|u-v\|^2 + \lambda TV(u)}{2\sigma^2}} du} \frac{\int u_j e^{-\frac{\|u-v\|^2 + \lambda TV(u)}{2\sigma^2}} du}{\int e^{-\frac{\|u-v\|^2 + \lambda TV(u)}{2\sigma^2}} du}.$$

Let us prove the continuity of dS_{LSE} again by dominated convergence theorem. Let i and j in Ω . For $g(u) = u_i u_j$ as for $g(u) = u_i$ or $g(u) = u_j$ or $g(u) = 1$, the functions $v \mapsto g(u) \exp(-\frac{\|u-v\|^2 + \lambda TV(u)}{2\sigma^2})$ are continuous with respect to v , and satisfy, for all h such that $\|h\| \leq \varepsilon$,

$$\left| g(u) e^{-\frac{\|u-v-h\|^2 + \lambda TV(u)}{2\sigma^2}} \right| \leq |g(u)| e^{-\frac{\frac{1}{2}\|u-v\|^2 + \lambda TV(u)}{2\sigma^2}} e^{\frac{\varepsilon^2}{2\sigma^2}},$$

which are integrable with respect to u , and now thanks to the dominated convergence theorem every term involved in dS_{LSE} is continuous, and thus dS_{LSE} is continuous. \square

3.1.4 Explicit solution for a 2-pixel clique

Contrary to MAP denoising where explicit solutions are computable for some kinds of images [Strong and Chan 1996, 2003; Steidl et al. 2004; Alter et al. 2005a], it seems much more difficult to find explicit solutions of TV-LSE, even for very simple images (except from constant images). Indeed, for an image as simple as a checkerboard or a white Dirac on one pixel, the LSE denoiser requires the integration over every image $u \in \mathbb{R}^\Omega$, and despite simplicity of the input and the multiple symmetry invariances the multi-dimensional integration is hard to handle.

A crucial parameter for explicit computation seems to be the size of the image. Explicit computation is feasible for very small images, which we will call *cliques* for this purpose, referring to the elementary interaction windows in Gibbs fields.

Let $v = (v_1, v_2)$ a 2-pixel clique. Then \hat{u}_{LSE} , here defined by

$$\hat{u}_{LSE} = \frac{\int (u_1, u_2) e^{-\frac{(u_1-v_1)^2+(u_2-v_2)^2+\lambda|u_1-u_2|}{2\sigma^2}} du_1 du_2}{\int e^{-\frac{(u_1-v_1)^2+(u_2-v_2)^2+\lambda|u_1-u_2|}{2\sigma^2}} du_1 du_2}, \quad (3.5)$$

with the particular *TV* operator

$$TV(u_1, u_2) = |u_1 - u_2|$$

can be explicitly derived, as described in the next proposition. The proof (which is quite a long proof for such a simple case) repels any generalization to large images.

Proposition 3.4 *Let v be a 2-pixel clique, with gray values (v_1, v_2) . Then the gray level of the right pixel of \hat{u}_{LSE} is given by*

$$\hat{u}_{LSE}(2) = v_2 - \frac{\lambda e^{\frac{-\lambda\delta}{2\sigma^2}} \phi\left(\frac{\delta-\lambda}{\sigma\sqrt{2}}\right) - e^{\frac{\lambda\delta}{2\sigma^2}} \phi\left(-\frac{\delta+\lambda}{\sigma\sqrt{2}}\right)}{2 e^{\frac{-\lambda\delta}{2\sigma^2}} \phi\left(\frac{\delta-\lambda}{\sigma\sqrt{2}}\right) + e^{\frac{\lambda\delta}{2\sigma^2}} \phi\left(-\frac{\delta+\lambda}{\sigma\sqrt{2}}\right)}, \quad (3.6)$$

where $\delta = v_2 - v_1$ and $\phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-\frac{t^2}{2}} dt$ is the cumulative distribution function of the normal distribution $\mathcal{N}(0, 1)$. The denoised left pixel $\hat{u}_{LSE}(1)$ can be deduced by symmetry.

The proof of Proposition 3.4 begins with the following lemma.

Lemma 3.1 *Let $J : \tau \mapsto \frac{1}{\sqrt{2\pi}} \int_{\mathbb{R}} e^{-\frac{(z-\tau)^2}{2}} \phi(z) dz$. Then for any $\tau \in \mathbb{R}$,*

$$J(\tau) = \phi\left(\frac{\tau}{\sqrt{2}}\right)$$

Proof of Lemma 3.1: J writes

$$J(\tau) = \frac{1}{2\pi} \int_z \int_{t < z} e^{-\frac{(z-\tau)^2}{2}} e^{-\frac{t^2}{2}} dt dz.$$

The 2-D change of variable $x = \frac{z-t}{\sqrt{2}}$ and $y = \frac{z+t}{\sqrt{2}}$ yields a separable integral

$$\begin{aligned} J(\tau) &= \frac{1}{2\pi} \int_y \int_{x>0} e^{-\frac{(\frac{x+y}{\sqrt{2}}-\tau)^2}{2}} e^{-\frac{(y-x)^2}{4}} dx dy \\ &= e^{-\frac{\tau^2}{2}} \frac{1}{2\pi} \int_y e^{-\frac{y^2-\sqrt{2}\tau y}{2}} dy \int_{x>0} e^{-\frac{x^2-\sqrt{2}\tau x}{2}} dx \\ &= \frac{e^{-\frac{\tau^2}{2}}}{2\pi} \cdot e^{\frac{\tau^2}{4}} \sqrt{2\pi} \cdot e^{\frac{\tau^2}{4}} \sqrt{2\pi} \phi\left(\frac{\tau}{\sqrt{2}}\right) = \phi\left(\frac{\tau}{\sqrt{2}}\right), \end{aligned}$$

which proves Lemma 3.1. \square

Proof of Proposition 3.4: The denominator of $\hat{u}_{LSE}(2)$ in (3.5)

$$D = \int_{\mathbb{R}^2} \exp\left[-\frac{1}{2\sigma^2} \left((u_1 - v_1)^2 + (u_2 - v_2)^2 + \lambda|u_1 - u_2|\right)\right] du_1 du_2$$

can be first integrated against $u_2 \in \mathbb{R}$, separating the cases $\{u_2 < u_1\}$ from $\{u_2 > u_1\}$. We have

$$D = \sigma \int_{u_1} e^{-\frac{(u_1-v_1)^2}{2\sigma^2}} \left[\phi\left(\frac{u_1 - v_2 + \frac{\lambda}{2}}{\sigma}\right) e^{\frac{(v_2 - \frac{\lambda}{2})^2 - \lambda u_1}{2\sigma^2}} + \phi\left(\frac{-u_1 + v_2 + \frac{\lambda}{2}}{\sigma}\right) e^{\frac{(v_2 + \frac{\lambda}{2})^2 + \lambda u_1}{2\sigma^2}} \right] du_1.$$

Let us make the change of variable $z = \frac{u_1 - v_2 + \frac{\lambda}{2}}{\sigma}$ in the first term and $z = \frac{u_1 - v_2 - \frac{\lambda}{2}}{\sigma}$ in the second one. We obtain

$$\begin{aligned} D &= \sqrt{2\pi}\sigma^2 e^{\frac{\lambda^2}{4\sigma^2}} \left[e^{\frac{\lambda\delta}{2\sigma^2}} J\left(-\frac{\delta + \lambda}{\sigma}\right) + e^{\frac{-\lambda\delta}{2\sigma^2}} J\left(\frac{\delta - \lambda}{\sigma}\right) \right] \\ &= 2\pi\sigma^2 e^{\frac{\lambda^2}{4\sigma^2}} \left[e^{\frac{\lambda\delta}{2\sigma^2}} \phi\left(-\frac{\delta + \lambda}{\sigma\sqrt{2}}\right) + e^{\frac{-\lambda\delta}{2\sigma^2}} \phi\left(\frac{\delta - \lambda}{\sigma\sqrt{2}}\right) \right], \end{aligned}$$

thanks to Lemma 3.1.

Now, let us focus on the computation of the numerator N_2 of $\hat{u}_{LSE}(2)$ in (3.5), that is on

$$N_2 = \int u_2 e^{-\frac{(u_1-v_1)^2 + (u_2-v_2)^2 + \lambda|u_1-u_2|}{2\sigma^2}} du_1 du_2.$$

We would like to link N_2 to the partial differential of D with respect to v_2 . To this end, let us first check that the theorem of differentiation under the integral sign applies for $\frac{\partial D}{\partial v_2}$. It suffices to check that for any compact interval $I \subset \mathbb{R}$, the function

$$|v_2 - u_2| e^{-\frac{(u_1-v_1)^2 + (u_2-v_2)^2 + \lambda|u_1-u_2|}{2\sigma^2}}$$

can be bounded by an integrable function independently from $v_2 \in I$. Let $I = [a, b]$. Then for any $v_2 \in [a, b]$, we have, thanks to a triangle inequality,

$$\forall u_2 \in \mathbb{R}, \quad \left| u_2 - \frac{a+b}{2} \right| - \frac{b-a}{2} \leq |u_2 - v_2| \leq \left| u_2 - \frac{a+b}{2} \right| + \frac{b-a}{2},$$

which implies that

$$|v_2 - u_2| e^{-\frac{(u_1 - v_1)^2 + (u_2 - v_2)^2 + \lambda |u_1 - u_2|}{2\sigma^2}} \leq \left(\left| u_2 - \frac{a+b}{2} \right| + \frac{b-a}{2} \right) e^{-\frac{(|u_2 - \frac{a+b}{2}| - \frac{b-a}{2})^2}{2\sigma^2}} e^{-\frac{(u_1 - v_1)^2}{2\sigma^2}},$$

which is integrable with respect to $(u_1, u_2) \in \mathbb{R}^2$. Then D is differentiable with respect to v_2 , and the partial differential of D with respect to v_2 satisfies

$$\frac{\partial D}{\partial v_2} = - \int \frac{v_2 - u_2}{\sigma^2} e^{-\frac{(u_1 - v_1)^2 + (u_2 - v_2)^2 + \lambda |u_1 - u_2|}{2\sigma^2}} du_1 du_2 = -\frac{1}{\sigma^2} [Dv_2 - N_2],$$

which implies that it is enough to compute $\frac{\partial D}{\partial v_2}$ to obtain

$$\hat{u}_{LSE}(2) = \frac{N_2}{D} = v_2 + \sigma^2 \frac{1}{D} \frac{\partial D}{\partial v_2}. \quad (3.7)$$

Now, we have

$$\frac{\partial D}{\partial v_2} = \frac{2\pi\lambda\sigma^2}{2\sigma^2} e^{\frac{\lambda^2}{2\sigma^2}} \left[e^{\frac{\lambda(v_2 - v_1)}{2\sigma^2}} \phi\left(-\frac{v_2 - v_1 + \lambda}{\sigma\sqrt{2}}\right) - e^{-\frac{\lambda(v_2 - v_1)}{2\sigma^2}} \phi\left(-\frac{v_2 - v_1 - \lambda}{\sigma\sqrt{2}}\right) \right],$$

which implies, by denoting $\delta = v_2 - v_1$, and by putting this together with (3.7), the desired equality (3.6). \square

Denoting $\bar{v} = \frac{v_1 + v_2}{2}$, and using the classical equivalent

$$\phi(x) \underset{x \rightarrow -\infty}{\sim} \frac{1}{|x|\sqrt{2\pi}} e^{-\frac{x^2}{2}},$$

we can find an equivalent for $\hat{u}_{LSE}(2)$ when $|\delta|$ is large

$$\begin{cases} \hat{u}_{LSE}(2) \sim \bar{v} + \frac{\delta}{2} - \frac{\lambda}{2} & \text{when } \delta \rightarrow +\infty, \\ \hat{u}_{LSE}(2) \sim \bar{v} + \frac{\delta}{2} + \frac{\lambda}{2} & \text{when } \delta \rightarrow -\infty. \end{cases}$$

The LSE denoising of 2-pixel clique should be compared to MAP denoising, for which the right gray level satisfies

$$\hat{u}_{MAP}(2) = \begin{cases} \bar{v} + \frac{\delta}{2} - \frac{\lambda}{2} & \text{if } \delta > \lambda \\ \bar{v} & \text{if } |\delta| \leq \lambda \\ \bar{v} + \frac{\delta}{2} + \frac{\lambda}{2} & \text{if } \delta < -\lambda. \end{cases}$$

We find out that when $|\delta| > \lambda$, the asymptotes of \hat{u}_{LSE} are exactly the affine lines of \hat{u}_{MAP} . This asymptotic behavior as well as the regularity of S_{LSE} (Proposition 3.3) can be seen on Figure 3.1, where $\hat{u}_{LSE}(2)$ is plotted against δ .




Figure 3.1: The LSE-denoised right pixel of a 2-pixel clique with gray levels $(-\delta/2, \delta/2)$ is plotted against δ (solid bold line) (see (3.6) with $\bar{v} = (v_1 + v_2)/2 = 0$, $\lambda = 20$, $\sigma = 10$), as well as its asymptotes (dashed lines). The MAP-denoised right pixel is also plotted against δ (solid thin line) and has the shape of a typical soft-thresholding. The outer branches of the MAP denoising are equal to the asymptotes of the LSE denoising. The LSE curve is more regular than the MAP curve, which agrees with the differentiability of S_{LSE} .

3.2 Asymptotic behavior of TV-LSE

Unlike MAP denoising, LSE denoising makes use of 2 distinct parameters λ and σ . If σ is fixed, we expect the regularity of the LSE denoised image to increase when λ increases as in MAP denoising. The impact of σ is less clear. The next theorem sums up several asymptotic behaviors of \hat{u}_{LSE} , when one of the parameter is fixed and the other goes to 0 or $+\infty$. The operator TV will only be assumed to be non-negative, except from the following cases for which we need other assumptions.

Assumption 1 (λ is fixed and $\sigma \rightarrow 0$) The TV operator is assumed to be convex and coercive.

Assumption 2 (λ is fixed and $\sigma \rightarrow \infty$) The TV operator is assumed to be Lipschitz continuous, i.e. there exists $\alpha > 0$ such that

$$\forall u, u' \in \mathbb{R}^\Omega, |TV(u) - TV(u')| \leq \alpha \|u' - u\|.$$

Assumption 3 (σ is fixed and $\lambda \rightarrow +\infty$) The TV operator is assumed to be defined on a finite rectangular grid $\Omega = \{1, \dots, n_1\} \times \{1, \dots, n_2\}$, and to satisfy

$$\exists C > 0 \quad / \quad \forall u \in \mathbb{R}^\Omega, \forall k, k' \in \Omega, |k - k'| = 1 \Rightarrow TV(u) \geq C |u_{k'} - u_k|.$$

Furthermore TV is assumed to be shift-invariant, i.e.

$$\forall c \in \mathbb{R}, \forall u \in \mathbb{R}^\Omega, TV(u + c) = TV(u).$$

Theorem 3.1 *Let v the noisy image. For all fixed $\lambda > 0$, we have*

$$\begin{aligned} (i) \quad \hat{u}_{LSE}(\lambda, \sigma) &\xrightarrow{\sigma \rightarrow 0} \hat{u}_{MAP}(\lambda), \\ (ii) \quad \hat{u}_{LSE}(\lambda, \sigma) &\xrightarrow{\sigma \rightarrow +\infty} v, \end{aligned}$$

and for all $\sigma > 0$, we have

$$\begin{aligned} (iii) \quad \hat{u}_{LSE}(\lambda, \sigma) &\xrightarrow{\lambda \rightarrow 0} v, \\ (iv) \quad \hat{u}_{LSE}(\lambda, \sigma) &\xrightarrow{\lambda \rightarrow +\infty} \bar{v}\mathbf{1}, \end{aligned}$$

where $\bar{v}\mathbf{1}$ is the image which is constant and equal to the average of v . Assumption 1 is necessary for item (i) only, Assumption 2 for item (ii) only, Assumption 3 for item (iv) only.

We summarize these limits, together with some others, in Figure 3.2.

Proof : When σ goes to 0, thanks to Assumption 1, the probability distribution $\frac{1}{Z} \exp\left(-\frac{E_\lambda}{2\sigma^2}\right)$ is log-concave, hence unimodal (Z is a normalizing constant depending on σ). Therefore it converges to the Dirac distribution in $\hat{u}_{MAP}(\lambda) = \arg \min_u E_\lambda(u)$, whose expectation is $\hat{u}_{MAP}(\lambda)$, which proves (i).

For (ii), let us consider the change of variable $w = (u - v)/\sigma$. Then

$$\hat{u}_{LSE}(\lambda, \sigma) = v + \frac{\int_{\mathbb{R}^\Omega} \sigma w e^{-\frac{1}{2}(\|w\|^2 + \frac{\lambda}{\sigma} TV(w + \frac{v}{\sigma}))} dw}{\int_{\mathbb{R}^\Omega} e^{-\frac{1}{2}(\|w\|^2 + \frac{\lambda}{\sigma} TV(w + \frac{v}{\sigma}))} dw} = v + \frac{N}{D}.$$

When $\sigma \rightarrow \infty$, the function inside the denominator D converges almost everywhere (a.e.) to $e^{-\|w\|^2/2}$, and is uniformly bounded by $e^{-\|w\|^2/2}$, thus thanks to Lebesgue's dominated convergence theorem, D converges towards $\int e^{-\|w\|^2/2} dw$.

For the numerator, notice that the mean value theorem applied to $x \mapsto e^{-x}$ implies the existence of a real number $c_{w,\sigma} \in [0, \frac{\lambda}{2\sigma} TV(w + \frac{v}{\sigma})]$ such that

$$e^{-\frac{\lambda}{2\sigma} TV(w + \frac{v}{\sigma})} = 1 - \frac{\lambda}{2\sigma} TV(w + \frac{v}{\sigma}) e^{-c_{w,\sigma}}.$$

Hence N can be split into

$$N = \sigma \int w e^{-\frac{\|w\|^2}{2}} dw - \frac{\lambda}{2} \int \underbrace{w e^{-\frac{\|w\|^2}{2}} TV(w + \frac{v}{\sigma}) e^{-c_{w,\sigma}}}_{f_\sigma(w)} dw.$$

The first integral is equal to zero. Concerning the second integral, when $\sigma \rightarrow \infty$, $c_{w,\sigma}$ goes to 0, and f_σ satisfies for every $\sigma \geq 1$,

$$\begin{cases} f_\sigma(w) &\xrightarrow{\sigma \rightarrow \infty} w e^{-\frac{\|w\|^2}{2}} TV(w) \quad \text{a.e.} \\ \|f_\sigma(w)\| &\leq \|w\| e^{-\frac{\|w\|^2}{2}} (TV(w) + \alpha \|v\|) \in L^1(\mathbb{R}^\Omega) \end{cases}$$

$\hat{u}_{LSE}(\lambda, \sigma) \longrightarrow \dots$					
when...	$\sigma \rightarrow 0$		σ fixed	$\sigma \rightarrow \infty$	
$\lambda \rightarrow 0$	v if $\frac{\lambda}{\sigma^2} \rightarrow \infty$	v if $\frac{\lambda}{\sigma^2} \rightarrow 0$	v	v	
λ fixed	$\hat{u}_{MAP}(\lambda)$		$\hat{u}_{LSE}(\lambda, \sigma)$	v	
$\lambda \rightarrow \infty$	$\bar{v}\mathbf{1}$		$\bar{v}\mathbf{1}$	$\bar{v}\mathbf{1}$ if $\frac{\lambda}{\sigma^2} \rightarrow \infty$	v if $\frac{\lambda}{\sigma^2} \rightarrow 0$

$\hat{u}_{LSE}(2\sigma^2\beta, \sigma) \longrightarrow \dots$					
when...	$\sigma \rightarrow 0$		σ fixed	$\sigma \rightarrow \infty$	
$\beta \rightarrow 0$	v		v	v if $\sigma^2\beta \rightarrow 0$	$\bar{v}\mathbf{1}$ if $\sigma^2\beta \rightarrow \infty$
β fixed	v		$\hat{u}_{LSE}(2\sigma^2\beta, \sigma)$	$\bar{v}\mathbf{1}$	
$\beta \rightarrow \infty$	v if $\sigma^2\beta \rightarrow 0$	$\bar{v}\mathbf{1}$ if $\sigma^2\beta \rightarrow \infty$	$\bar{v}\mathbf{1}$	$\bar{v}\mathbf{1}$	

Figure 3.2: Top table: behavior of $\hat{u}_{LSE}(\lambda, \sigma)$ for extreme values of λ and σ . Notice that the case where λ and σ simultaneously go to 0 or ∞ splits into several subcases, each concerning a behavior of $\frac{\lambda}{2\sigma^2} = \beta$ (see Subsection 1.2.1 for the interpretation of β as the inverse of the temperature coefficient in the Gibbs field given by the prior $\frac{1}{Z}e^{-\beta TV(u)}$). In these cases we only give the limit of \hat{u}_{LSE} when $\beta \rightarrow 0$ and $\beta \rightarrow \infty$. Bottom table: parameters (β, σ) are also suitable parameters for \hat{u}_{LSE} . Letting the parameters fixed or go to 0 or ∞ leads to the second table. A dichotomy is also needed in the cases $(\beta \rightarrow 0, \sigma \rightarrow \infty)$ and $(\beta \rightarrow \infty, \sigma \rightarrow 0)$, between the cases $\lambda \rightarrow 0$ and $\lambda \rightarrow \infty$. These tables will be experimentally confirmed in Chapter 4 Section 4.2.

where α is the Lipschitz-continuity coefficient of TV (Assumption 2). Again Lebesgue's dominated convergence theorem applies, and

$$\int f_\sigma(w)dw \xrightarrow{\sigma \rightarrow \infty} \int we^{-\frac{\|w\|^2}{2}} TV(w)dw = 0$$

because the function inside the integral is odd, and then N goes to 0, which implies the convergence of $\hat{u}_{LSE}(\lambda, \sigma)$ towards v , and proves (ii).

For (iii), the numerator and the denominator of $\hat{u}_{LSE}(k)$ satisfy

$$u_k e^{-\frac{\|u-v\|^2 + \lambda TV(u)}{2\sigma^2}} \xrightarrow{\lambda \rightarrow 0} u_k e^{-\frac{\|u-v\|^2}{2\sigma^2}} \quad \text{and} \quad e^{-\frac{\|u-v\|^2 + \lambda TV(u)}{2\sigma^2}} \xrightarrow{\lambda \rightarrow 0} e^{-\frac{\|u-v\|^2}{2\sigma^2}}$$

and the uniform bounds

$$\left| u_k e^{-\frac{\|u-v\|^2 + \lambda TV(u)}{2\sigma^2}} \right| \leq |u_k| e^{-\frac{\|u-v\|^2}{2\sigma^2}} \in L^1(\mathbb{R}^\Omega) \quad \text{and} \quad \left| e^{-\frac{\|u-v\|^2 + \lambda TV(u)}{2\sigma^2}} \right| \leq e^{-\frac{\|u-v\|^2}{2\sigma^2}} \in L^1(\mathbb{R}^\Omega)$$

give the result thanks again to Lebesgue's dominated convergence theorem.

For (iv), the dominated convergence theorem cannot be directly applied because $ue^{-\frac{1}{2\sigma^2}TV(u)}$ is not $L^1(\mathbb{R}^\Omega)$. We need to come down to a space of constant mean images where it becomes L^1 . Let \bar{v} denote the average of v on Ω . Then thanks to Assumption 3, the change of variable $x = u - \bar{v}$ yields

$$\hat{u}_{LSE}(\lambda, \sigma) = \bar{v} + \frac{\int x e^{-\frac{1}{2\sigma^2}(\|x+\bar{v}-v\|^2 + \lambda TV(x))} dx}{\int e^{-\frac{1}{2\sigma^2}(\|x+\bar{v}-v\|^2 + \lambda TV(x))} dx},$$

and denoting \mathcal{E}_0 the space of zero-mean images, splitting x into $x = \bar{x} + \varepsilon$ with $\varepsilon \in \mathcal{E}_0$ yields

$$\begin{aligned} \hat{u}_{LSE}(\lambda, \sigma) &= \bar{v} + \frac{\int_{\varepsilon \in \mathcal{E}_0} e^{-\frac{1}{2\sigma^2}(\|\varepsilon-v\|^2 + \lambda TV(\varepsilon))} \int_{\bar{x}} (\bar{x} + \varepsilon) e^{-\frac{1}{2\sigma^2}(\bar{x}^2 - \bar{v}^2)} d\bar{x} d\varepsilon}{\int_{\varepsilon \in \mathcal{E}_0} e^{-\frac{1}{2\sigma^2}(\|\varepsilon-v\|^2 + \lambda TV(\varepsilon))} \int_{\bar{x}} e^{-\frac{1}{2\sigma^2}(\bar{x}^2 - \bar{v}^2)} d\bar{x} d\varepsilon} \\ &= \bar{v} + \frac{\int_{\varepsilon \in \mathcal{E}_0} e^{-\frac{1}{2\sigma^2}(\|\varepsilon-v\|^2 + \lambda TV(\varepsilon))} (0 + \sqrt{2\pi}\sigma\varepsilon) e^{\frac{\bar{v}^2}{2\sigma^2}} d\varepsilon}{\int_{\varepsilon \in \mathcal{E}_0} e^{-\frac{1}{2\sigma^2}(\|\varepsilon-v\|^2 + \lambda TV(\varepsilon))} \sqrt{2\pi}\sigma e^{\frac{\bar{v}^2}{2\sigma^2}} d\varepsilon} \\ &= \bar{v} + \frac{\int_{\varepsilon \in \mathcal{E}_0} \varepsilon e^{-\frac{1}{2\sigma^2}(\|\varepsilon-v\|^2 + \lambda TV(\varepsilon))} d\varepsilon}{\int_{\varepsilon \in \mathcal{E}_0} e^{-\frac{1}{2\sigma^2}(\|\varepsilon-v\|^2 + \lambda TV(\varepsilon))} d\varepsilon}. \end{aligned}$$

Now the change of variable $y = \lambda\varepsilon$ leads to

$$\hat{u}_{LSE}(\lambda, \sigma) = \bar{v} + \frac{1}{\lambda} \frac{\int_{y \in \mathcal{E}_0} y e^{-\frac{1}{2\sigma^2}(\|\frac{y}{\lambda}-v\|^2 + TV(y))} dy}{\int_{y \in \mathcal{E}_0} e^{-\frac{1}{2\sigma^2}(\|\frac{y}{\lambda}-v\|^2 + TV(y))} dy}. \quad (3.8)$$

For both functions $g(y) = 1$ and $g(y) = y$, we have

$$\left\{ \begin{array}{l} g(y) e^{-\frac{1}{2\sigma^2}(\|\frac{y}{\lambda}-v\|^2 + TV(y))} \xrightarrow{\lambda \rightarrow \infty} g(y) e^{-\frac{1}{2\sigma^2}(\|v\|^2 + TV(y))}, \\ \left\| g(y) e^{-\frac{1}{2\sigma^2}(\|\frac{y}{\lambda}-v\|^2 + TV(y))} \right\| \leq \|g(y)\| e^{-\frac{1}{2\sigma^2}TV(y)} \text{ which is } L^1(\mathcal{E}_0) \text{ as we check now.} \end{array} \right.$$

First, for $g(y) = 1$, Assumption 3 implies the existence of a constant $C_1 > 0$ such that

$$TV(y) \geq C_1 \left[\sum_{i=1}^{n_1} \sum_{j=0}^{n_2-1} |y_{i,j+1} - y_{i,j}| + \sum_{j=1}^{n_2-1} |y_{i+1,1} - y_{i,1}| \right].$$

Thus we get

$$\left| e^{-\frac{1}{2\sigma^2}TV(y)} \right| \leq e^{-\frac{C_1}{2\sigma^2} \left[\sum_{i=1}^{n_1} \sum_{j=1}^{n_2-1} |y_{i,j+1} - y_{i,j}| + \sum_{j=1}^{n_2-1} |y_{i+1,1} - y_{i,1}| \right]},$$

and the change of variables $y \in \mathcal{E}_0 \mapsto ((\eta_{i,j})_{1 \leq i \leq n_1, 1 \leq j \leq n_2-1}, (\xi_i)_{1 \leq i \leq n_1-1})$ with $\eta_{i,j} = y_{i,j+1} - y_{i,j}$ and $\xi_i = y_{i+1,1} - y_{i,1}$ leads to

$$\left| \int e^{-\frac{1}{2\sigma^2}TV(y)} dy \right| \leq \prod_{i=1}^{n_1} \prod_{j=1}^{n_2-1} \int e^{-\frac{C_1}{2\sigma^2}|\eta_{i,j}|} d\eta_{i,j} \cdot \prod_{i=1}^{n_1-1} \int e^{-\frac{C_1}{2\sigma^2}|\xi_i|} d\xi_i$$

which is finite because $|\Omega|$ is finite. Now for $g(y) = y$, we have the equivalence of norms on \mathcal{E}_0

$$\|y\|_2 \asymp \|y\|_1 \asymp \sum_{i=1}^{n_1} \sum_{j=1}^{n_2-1} |y_{i,j+1} - y_{i,j}| + \sum_{j=1}^{n_2-1} |y_{i+1,1} - y_{i,1}|.$$

(The notation $f_1 \asymp f_2$ means that there exist two *positive* constants A and B such that $Af_1 \leq f_2 \leq Bf_1$). Then there exists $C_2 > 0$ such that

$$\forall y \in \mathcal{E}_0, \quad \|y\|_2 \leq C_2 \left[\sum_{i=1}^{n_1} \sum_{j=1}^{n_2-1} |y_{i,j+1} - y_{i,j}| + \sum_{j=1}^{n_2-1} |y_{i+1,1} - y_{i,1}| \right].$$

Now, with the same change of variables as in the case $g(y) = 1$, i.e.

$$y \in \mathcal{E}_0 \mapsto ((\eta_{i,j})_{1 \leq i \leq n_1, 1 \leq j \leq n_2-1}, (\xi_i)_{1 \leq i \leq n_1-1}),$$

we get

$$\begin{aligned} & \int_{\mathcal{E}_0} \|y\|_2 e^{-\frac{1}{2\sigma^2}TV(y)} dy \\ & \leq C_2 \int_{\mathcal{E}_0} \left[\sum_{i=1}^{n_1} \sum_{j=1}^{n_2-1} |y_{i,j+1} - y_{i,j}| + \sum_{j=1}^{n_2-1} |y_{i+1,1} - y_{i,1}| \right] e^{-\frac{1}{2\sigma^2}TV(y)} dy \\ & \leq C_2 \left[\sum_{i=1}^{n_1} \sum_{j=1}^{n_2-1} \int |\eta_{i,j}| e^{-\frac{C_1}{2\sigma^2} [\sum_{i'=1}^{n_1} \sum_{j'=1}^{n_2-1} |\eta_{i',j'}| + \sum_{i'=1}^{n_1-1} |\xi_{i'}|]} d\eta_{i,j} d\xi_i \right. \\ & \quad \left. + \sum_{j=1}^{n_2-1} \int |\xi_i| e^{-\frac{C_1}{2\sigma^2} [\sum_{i'=1}^{n_1} \sum_{j'=1}^{n_2-1} |\eta_{i',j'}| + \sum_{i'=1}^{n_1-1} |\xi_{i'}|]} d\eta_{i,j} d\xi_i \right] \\ & = C_2 \left[\sum_{i=1}^{n_1} \sum_{j=1}^{n_2-1} \int |\eta_{i,j}| e^{-\frac{C_1|\eta_{i,j}|}{2\sigma^2}} d\eta_{i,j} \prod_{\{i',j'\} \neq \{i,j\}} \int e^{-\frac{C_1|\eta_{i',j'}|}{2\sigma^2}} d\eta_{i',j'} \prod_{i'=1}^{n_1-1} \int e^{-\frac{C_1|\xi_{i'}|}{2\sigma^2}} d\xi_{i'} \right. \\ & \quad \left. + \sum_{i=1}^{n_1-1} \prod_{i'=1}^{n_1} \prod_{j'=0}^{n_2-1} \int e^{-\frac{C_1|\eta_{i',j'}|}{2\sigma^2}} d\eta_{i',j'} \int |\xi_i| e^{-\frac{C_1|\xi_i|}{2\sigma^2}} d\xi_i \prod_{i' \neq i} \int e^{-\frac{C_1|\xi_{i'}|}{2\sigma^2}} d\xi_{i'} \right] \end{aligned}$$

which is finite. Then thanks again to Lebesgue's dominated convergence theorem, the fraction of integrals in (3.8) is bounded uniformly in λ and dividing by λ leads to the desired convergence

$$\hat{u}_{LSE}(\lambda, \sigma) \xrightarrow{\lambda \rightarrow \infty} \bar{v}.$$

□

3.3 No staircasing in TV-LSE denoised images

Now we come to the main improvement of TV-LSE denoising compared to TV-MAP denoising: there is no staircasing in TV-LSE denoised images. Here again, the domain Ω is a finite subset of \mathbb{Z}^2 , and does not need to be a rectangle. The noisy input image is assumed to be a random image, which will be denoted V . We need the two following assumptions.

Assumptions

(A1) The distribution of the random image V is absolutely continuous with respect to Lebesgue's measure on \mathbb{R}^Ω .

(A2) The TV operator is any mapping $\mathbb{R}^\Omega \rightarrow \mathbb{R}$ such that, for every image u ,

$$\begin{cases} \forall c \in \mathbb{R}, \forall u \in \mathbb{R}^\Omega, & TV(u + c) = TV(u) \\ \exists \alpha > 0 & / \quad \forall u \in \mathbb{R}^\Omega, \forall h \in \mathbb{R}, \forall k \in \Omega, & TV(u) - \alpha|h| \leq TV(u^{k,h}) \leq TV(u) + \alpha|h| \\ \exists C > 0 & / \quad \forall k, k' \in \Omega, \quad |k - k'| = 1 \Rightarrow & TV(u) \geq C|u_{k'} - u_k| \end{cases}$$

where $u^{k,h}$ is equal to u except at pixel k where the value is $u_k^{k,h} = u_k + h$.

We now come to the statement of the main theorem.

Theorem 3.2 *Let V be a random image satisfying Assumption (A1). Let $S_{LSE}(V)$ denote the output of V through the TV-LSE denoiser, where the TV operator satisfies (A2). Let k and k' be neighbor pixels of Ω , i.e. such that $|k' - k| = 1$. Then $S_{LSE}(V)$ satisfies*

$$\mathbb{P}(S_{LSE}(V)_{k'} = S_{LSE}(V)_k) = 0.$$

In other words, for almost every realization v of V , \hat{u}_{LSE} contains no region with 2 pixels or more where it is constant: there is no staircasing, in the sense of [Nikolova 2004] (no constant region, but perhaps discontinuities, see Chapter 1).

Examples and interpretation

- If V writes $V = u + N$ with u a fixed image and N a white Gaussian noise, i.e. a realization of V is a noisy image, then V satisfies Assumption (A1), and thanks to the Theorem, \hat{u}_{LSE} almost surely contains no staircasing.
- If V is drawn from the total variation distribution (i.e. $V \sim \frac{1}{Z} e^{-\lambda TV(V)}$), then it satisfies Assumption (A1), and \hat{u}_{LSE} almost surely contains no staircasing.
- Most of the commonly used TV operators satisfy Assumption (A2), as Proposition 3.5 will precise at Section 3.4.

Proof of Theorem 3.2: Let $v \in \mathbb{R}^\Omega$ be an input image. Let $g : \mathbb{R} \rightarrow \mathbb{R}$ the function defined by

$$g : z \mapsto \int_{\mathbb{R}^\Omega} (u_{k'} - u_k) \exp \left(-\frac{(u_k - z)^2 + \sum_{l \neq k} (u_l - v_l)^2 + \lambda TV(u)}{2\sigma^2} \right) du. \quad (3.9)$$

We have

$$S_{LSE}(v)_{k'} = S_{LSE}(v)_k \iff g(v_k) = 0.$$

We need the following lemma which we shall prove later.

Lemma 3.2 *g is analytic and non-identically null.*

Assume Lemma 3.2 proven. Then $g^{-1}(\{0\})$ cannot contain any accumulation point (otherwise g would be constant thanks to the isolated zero theorem for analytic functions), and in particular the event $(g(z) = 0)$ has a zero probability under the density distribution of the gray level V_k of V because of Assumption **(A1)**

$$\mathbb{P}_{V_k}(g(V_k) = 0) = 0.$$

Now let f be the probability density function of V . We have

$$\mathbb{P}(g(V_k) = 0) = \int_{(v_l)_{l \neq k}} \int_{v_k} \mathbf{1}_{g(v_k)=0} f(v) dv_k d(v_l)_{l \neq k}$$

with $\mathbf{1}_{g(v_k)=0} = 0$ almost everywhere. The notation $d(v_l)_{l \neq k}$ means an integration over all the gray levels of v , except from v_k . Therefore

$$\mathbb{P}(g(V_k) = 0) = \int_{(v_l)_{l \neq k}} \int_{v_k} 0 \cdot f(v) dv_k d(v_l)_{l \neq k} = 0.$$

Finally we get

$$\mathbb{P}(S_{LSE}(V)(k) = S_{LSE}(V)(k')) = \mathbb{P}(g(V_k) = 0) = 0,$$

which ends the proof. \square

Proof of Lemma 3.2: Let $z \in \mathbb{C}$ generalize the gray level v_k to complex values. With (3.9), g can naturally be extended to a $\mathbb{C} \rightarrow \mathbb{C}$ function. Let us prove that g is holomorphic. To this end, let $f : \mathbb{R}^\Omega \times \mathbb{C} \rightarrow \mathbb{C}$ denote the function inside the integral defining g . We have

$$f(u, z) = (u_{k'} - u_k) \exp \left(-\frac{(u_k - z)^2 + \sum_{l \neq k} (u_l - v_l)^2 + \lambda TV(u)}{2\sigma^2} \right).$$

For any $A > 0$, let K_A denote the compact set $[-A, A] + i[-A, A]$.

- for all $z \in K_A$, $u \mapsto f(u, z)$ is integrable;
- for every $u \in \mathbb{R}^\Omega$, $z \mapsto f(u, z)$ is Gaussian, hence holomorphic;

- for every $z = x + iy \in K_A$ ($x, y \in \mathbb{R}$), f can be bounded from above as follows

$$\begin{aligned}
|f(u, z)| &\leq |u_{k'} - u_k| \left| \exp \left(-\frac{\sum_{l \neq k} (u_l - v_l)^2 + (u_k - z)^2}{2\sigma^2} \right) \right| \\
&= |u_{k'} - u_k| \exp \left(-\frac{\sum_{l \neq k} (u_l - v_l)^2 + (u_k - x)^2 - y^2}{2\sigma^2} \right) \\
&\leq |u_{k'} - u_k| \exp \left(-\frac{\sum_{l \neq k} (u_l - v_l)^2 + (u_k^2 - 2A|u_k| - A^2)}{2\sigma^2} \right) \\
&\leq (|u_{k'}| + |u_k|) \exp \left(-\frac{\sum_{l \neq k} (u_l - v_l)^2 + (|u_k| - A)^2}{2\sigma^2} \right) e^{\frac{A^2}{\sigma^2}},
\end{aligned}$$

which is integrable with respect to $u \in \mathbb{R}^\Omega$.

Thus, applying the theorem of differentiation of holomorphic functions under Lebesgue integral, $g : z \mapsto \int f(u, z) du$ is holomorphic on K_A , for every A , thus is holomorphic on \mathbb{C} . It comes that g is an analytic function on \mathbb{C} .

Now we must prove that g is not the constant null function. Let us assume that g is null everywhere. Then

$$\begin{aligned}
g(z) &= \int_u (u_{k'} - u_k) \exp \left(-\frac{(u_k - z)^2 + \sum_{l \neq k} (u_l - v_l)^2 + \lambda TV(u)}{2\sigma^2} \right) du \\
&= \int_{u_k} e^{-\frac{(z - u_k)^2}{2\sigma^2}} \varphi(u_k) du_k \\
&= G_\sigma * \varphi(z)
\end{aligned}$$

where G_σ is the Gaussian kernel with bandwidth σ , and φ is given by

$$\varphi(u_k) = \int_{(u_l)_{l \neq k}} (u_{k'} - u_k) \exp \left(-\frac{\sum_{l \neq k} (u_l - v_l)^2 + \lambda TV(u)}{2\sigma^2} \right) d(u_l)_{l \neq k} \quad (3.10)$$

Let us first prove that φ is identically null. To this end, let us check that φ is in L^1 , using Assumption **(A2)**. We have

$$\begin{aligned}
\int |\varphi(u_k)| du_k &= \int_{\mathbb{R}^\Omega} |u_{k'} - u_k| \exp \left(-\frac{\sum_{l \neq k} (u_l - v_l)^2 + \lambda TV(u)}{2\sigma^2} \right) du \\
&\leq \int_{\mathbb{R}^\Omega} |u_{k'} - u_k| \exp \left(-\frac{\sum_{l \neq k} (u_l - v_l)^2 + \lambda C |u_{k'} - u_k|}{2\sigma^2} \right) du \\
&\leq \int_{\mathbb{R}} |z| \exp \left(-\frac{C|z|}{2\sigma^2} \right) dz \cdot \prod_{l \neq k} \int_{\mathbb{R}} \exp \left(-\frac{(u_l - v_l)^2}{2\sigma^2} \right) du_l \\
&< +\infty.
\end{aligned}$$

Hence φ is in L^1 and satisfies

$$\forall z \in \mathbb{R} \quad G_\sigma * \varphi(z) = 0.$$

Considering Fourier transforms (written $\hat{\cdot}$) we get

$$\forall \xi \in \mathbb{R} \quad \hat{G}_\sigma(\xi) \cdot \hat{\varphi}(\xi) = 0.$$

As $\hat{G}_\sigma(z) = G_{2\pi/\sigma}$ never vanishes, $\hat{\varphi}$ necessarily is identically null. Injectivity of Fourier transform implies that φ is identically null.

Now we shall find a contradiction by proving that $\varphi(z)$ is negative for a large enough $z \in \mathbb{R}$. The change of variable $u_l \leftarrow u_l - z$ for every $l \neq k$ in Equation (3.10) yields, using Assumption **(A2)**,

$$\varphi(z) = \int_{(u_l)_{l \neq k}} u_{k'} \exp\left(-\frac{\sum_{l \neq k} (u_l + z - v_l)^2 + \lambda TV(u)}{2\sigma^2}\right) d(u_l)_{l \neq k}$$

where u in $TV(u)$ is composed of the gray levels u_l if $l \neq k$ and 0 if $l = k$. Let $h = u_{k'}$. Then

$$\varphi(z) = \int_{(u_l)_{l \neq k, k'}} \int_{h \in \mathbb{R}} h \exp\left(-\frac{\sum_{l \neq k, k'} (u_l + z - v_l)^2 + (h + z - v_{k'})^2 + \lambda TV(u^{k', h})}{2\sigma^2}\right) dh d(u_l)_{l \neq k, k'},$$

where $u^{k', h}$ is defined by

$$u_l^{k', h} = \begin{cases} 0 & \text{if } l = k \\ h & \text{if } l = k' \\ u_l & \text{otherwise.} \end{cases}$$

Now, coupling positive and negative values of h in the integral, we get

$$\varphi(z) = \int_{(u_l)_{l \neq k, k'}} \exp\left(-\frac{\sum_{l \neq k, k'} (u_l + z - v_l)^2}{2\sigma^2}\right) \int_{h > 0} h \phi_z(h, (u_l)_{l \neq k, k'}) dh d(u_l)_{l \neq k, k'} \quad (3.11)$$

with

$$\phi_z(h, (u_l)_{l \neq k, k'}) = \exp\left(-\frac{(h + z - v_{k'})^2 + \lambda TV(u^{k', h})}{2\sigma^2}\right) - \exp\left(-\frac{(-h + z - v_{k'})^2 + \lambda TV(u^{k', -h})}{2\sigma^2}\right).$$

Let us prove the strict negativity of ϕ_z when z is large enough. The assumption **(A2)** is needed to bound ϕ_z above by 0. We have

$$\phi_z(h, (u_l)_{l \neq k, k'}) \leq \exp\left(-\frac{h^2 + (z - v_{k'})^2 + \lambda TV(u^{k', 0})}{2\sigma^2}\right) \cdot A,$$

with

$$\begin{aligned} A &= \exp\left(\frac{-2h(z - v_{k'}) + \lambda \alpha h}{2\sigma^2}\right) - \exp\left(\frac{2h(z - v_{k'}) - \lambda \alpha h}{2\sigma^2}\right) \\ &= (-2) \sinh\left(\frac{h(z - v_{k'} - \frac{\lambda \alpha}{2})}{\sigma^2}\right). \end{aligned}$$

Assume $z > v_{k'} + \frac{\lambda\alpha}{2}$. Then $A < 0$, and therefore $\phi_z < 0$. As $\varphi(z)$ is the integral of the product of negative functions on a non-negligible set (3.11), it is negative. This contradicts the nullity of $\varphi(z)$ for large enough $z \in \mathbb{R}$. Finally g is non identically null. This concludes the proof of Lemma 3.2. \square

The Assumption **(A2)** about the total variation is satisfied by the most common discretizations of TV . For example, TV operators defined by

$$TV(u) = \sum_{x,y} (|u(x+1,y) - u(x,y)| + |u(x,y+1) - u(x,y)|)$$

or

$$TV(u) = \sum_{x,y} \sqrt{|u(x+1,y) - u(x,y)|^2 + |u(x,y+1) - u(x,y)|^2}$$

satisfy Assumption **(A2)**. These schemes are particular cases of more general schemes, which are proven in Section 3.4, Proposition 3.5, to all satisfy Assumption **(A2)**, as well as all the assumptions which were needed throughout the thesis.

3.4 Brief analysis of the assumptions on discrete TV

In this section we briefly compile all the assumptions that have been made on the TV operator both in Chapter 2 and in this chapter, and prove that they are satisfied by the most commonly used discretizations of the continuous total variation in finite dimension.

Proposition 3.5 *Let \mathcal{N} be a neighborhood shape, i.e. a finite subset of \mathbb{Z}^2 containing 0, such that $(0,1) \in \mathcal{N}$ and $(1,0) \in \mathcal{N}$. Let $|\mathcal{N}|$ denote the cardinality of \mathcal{N} . For each $k \in \Omega$ let \mathcal{N}_k denote the neighborhood \mathcal{N} centered at pixel k defined by*

$$\mathcal{N}_k = k + \mathcal{N}.$$

$\dot{\mathcal{N}}_k$ will denote the enucleated neighborhood $\mathcal{N}_k \setminus \{k\}$. Let $1 \leq p \leq \infty$ and let ε be an image of \mathbb{R}^Ω satisfying

$$\forall k \in \mathbb{R}^\Omega, \quad \varepsilon_k \geq 0$$

(the case $\varepsilon = 0$ is possible indeed). If $A \subset \Omega$ and $u \in \mathbb{R}^\Omega$, $u(A)$ will denote the image u restricted to the set A . For all pixel $k \in \Omega$, let $\|\cdot\|_k$ denote any norm over $\dot{\mathcal{N}}_k \cap \Omega$. Let us define the TV operator by

$$TV_{\varepsilon,p}(u) = \sum_{k \in \Omega} \left(\varepsilon_k^p + \|u(\dot{\mathcal{N}}_k \cap \Omega) - u_k\|_k^p \right)^{1/p}$$

and let $\pi_{\varepsilon,p}(u) = \frac{1}{Z} e^{-\frac{\|u-v\|^2 + \lambda TV_{\varepsilon,p}(u)}{2\sigma^2}}$ be the associate posterior distribution. Then $TV_{\varepsilon,p}$ and $\pi_{\varepsilon,p}$

satisfy the following assertions

- (i) $\pi_{\varepsilon,p}$ admits finite second order moments
- (ii) $\pi_{\varepsilon,p}$ is absolutely continuous w.r.t. Lebesgue's measure with positive and continuous density on \mathbb{R}^Ω
- (iii) $\exists \alpha, \forall u, h \in \mathbb{R}^\Omega, |TV_{\varepsilon,p}(u+h) - TV_{\varepsilon,p}(u)| \leq \alpha \|h\|$
- (iv) $\forall s \in (0,1), \exists a > 0, \forall u \in \mathbb{R}^\Omega, \pi_{\varepsilon,p}^{-s}(u) \geq a \|u\|$
- (v) $\forall u \in \mathbb{R}^\Omega, TV_{\varepsilon,p}(u) \geq 0$
- (vi) $\forall c \in \mathbb{R}, \forall u \in \mathbb{R}^\Omega, TV_{\varepsilon,p}(u+c) = TV(u)$
- (vii) $TV_{\varepsilon,p}$ is convex
- (viii) $\exists C > 0, \forall u \in \mathbb{R}^\Omega, \forall k, k' \in \Omega, |k - k'| = 1 \Rightarrow TV_{\varepsilon,p}(u) \geq C |u_{k'} - u_k|$.

Items (i) – (iv) are required in the theoretical results of Chapter 2. Item (v) is a general assumption for all the result of this chapter. (iii), (vi), (vii) and (viii) are required for the proof of Theorem 3.1 about the asymptotic behavior of \hat{u}_{LSE} . Last, items (iii), (vi) and (viii) are required in the proof of Theorem 3.2 stating the absence of staircasing.

Proof : Items (ii), (iv), (v), (vi) are straightforward. We focus on the other items.

Now (i) holds because $TV_{\varepsilon,p}$ is non-negative.

To prove (vii), notice that for each $k \in \Omega$,

$$(\varepsilon_k, u(\dot{\mathcal{N}}_k \cap \Omega)) \mapsto \left(\varepsilon_k^p + \|u(\dot{\mathcal{N}}_k \cap \Omega)\|_k^p \right)^{1/p}$$

is a norm over $\mathbb{R} \times (\dot{\mathcal{N}}_k \cap \Omega)$, which will be denoted $\|\cdot\|_{k+}$. Therefore $\|\cdot\|_{k+}$ is convex for each $k \in \Omega$. Their restriction to $\{\varepsilon_k\} \times (\dot{\mathcal{N}}_k \cap \Omega)$ remains convex. $TV_{\varepsilon,p}$ is a positive combination of these convex functions, and hence is convex.

We use the norms $\|\cdot\|_{k+}$ to prove (iii). Let $u \in \mathbb{R}^\Omega$, $k \in \Omega$ and $h \in \mathbb{R}$. Let $u^{k,h}$ be the image equal to u except at pixel k where its value is $u_k^{k,h} = u_k + h$. Then there exists $c_k > 0$ independent of u such that

$$\begin{aligned} \|(\varepsilon_k, u^{k,h}(\dot{\mathcal{N}}_k \cap \Omega)) - u_k\|_{k+} &= \|(\varepsilon_k, u(\dot{\mathcal{N}}_k \cap \Omega)) - u_k - h\|_{k+} \\ &\leq \|(\varepsilon_k, u(\dot{\mathcal{N}}_k \cap \Omega)) - u_k\|_{k+} + \|(0, h \mathbb{1}_{\dot{\mathcal{N}}_k \cap \Omega})\|_{k+} \\ &\leq \|(\varepsilon_k, u(\dot{\mathcal{N}}_k \cap \Omega)) - u_k\|_{k+} + c_k |h| \end{aligned}$$

thanks to the triangle inequality and because $\|\cdot\|_{k+}$ is equivalent to $\|\cdot\|_1$ in finite dimension. Likewise

$$\|(\varepsilon_k, u^{k,h}(\dot{\mathcal{N}}_k \cap \Omega)) - u_k\|_{k+} \geq \|(\varepsilon_k, u(\dot{\mathcal{N}}_k \cap \Omega)) - u_k\|_{k+} - c_k |h|.$$

Now let $l \in \dot{\mathcal{N}}_k$. Then there exists $c_l > 0$ independent of u such that

$$\begin{aligned} \|(\varepsilon_l, u^{k,h}(\dot{\mathcal{N}}_l \cap \Omega)) - u_l\|_{l+} &\leq \|(\varepsilon_l, u(\dot{\mathcal{N}}_l \cap \Omega)) - u_l\|_{l+} + \|(0, 0^{k,h}(\dot{\mathcal{N}}_l \cap \Omega))\|_{l+} \\ &\leq \|(\varepsilon_l, u(\dot{\mathcal{N}}_l \cap \Omega)) - u_l\|_{l+} + c_l |h| \end{aligned}$$

and

$$\|(\varepsilon_l, u^{k,h}(\dot{\mathcal{N}}_l \cap \Omega)) - u_l\|_{l_+} \geq \|(\varepsilon_l, u(\dot{\mathcal{N}}_l \cap \Omega)) - u_l\|_{l_+} - c_l|h|$$

Thus, as $TV_{\varepsilon,p}(u^{k,h})$ satisfies

$$TV_{\varepsilon,p}(u^{k,h}) = \sum_{l \in \Omega} \|(\varepsilon_l, u^{k,h}(\dot{\mathcal{N}}_l \cap \Omega)) - u_l\|_{l_+},$$

gathering the obtained equalities, we get

$$TV_{\varepsilon,p}(u) - \left(\sum_{l \in \Omega} c_l\right)|h| \leq TV_{\varepsilon,p}(u^{k,h}) \leq TV_{\varepsilon,p}(u) + \left(\sum_{l \in \Omega} c_l\right)|h|$$

which proves (iii) since $\alpha = \sum_{l \in \Omega} c_l$ does not depend on u nor h .

To prove (viii), let k and k' such that $|k' - k| = 1$. Exchanging k and k' if necessary we can assume $k' \in \mathcal{N}_k$. Then there exists $C_k > 0$ depending on the norm $\|\cdot\|_{k_+}$ only such that

$$\begin{aligned} TV_{\varepsilon,p}(u) &\geq \|(\varepsilon_k, u^{k,h}(\dot{\mathcal{N}}_k \cap \Omega)) - u_k\|_{k_+} \\ &\geq C_k \|(\varepsilon_k, u^{k,h}(\dot{\mathcal{N}}_k \cap \Omega)) - u_k\|_1 \\ &\geq C_k |u_{k'} - u_k| \end{aligned}$$

because in finite dimension, $\|\cdot\|_{k_+}$ is equivalent to $\|\cdot\|_1$. Now letting

$$C = \min_{l \in \Omega} C_l,$$

we obtain a constant $C > 0$ independent of k such that

$$TV_{\varepsilon,p}(u) \geq C |u_{k'} - u_k|$$

holds, which proves (viii). \square

3.5 Global analysis of LSE denoisers built on a more general prior

In this section we take an analytic point of view on the TV-LSE denoiser. We consider general prior distributions satisfying basic requirements only, and we prove that LSE denoisers are invertible and regular, and satisfy stability properties, among other things.

Afterwards, the TV-LSE denoiser is proven to be a Maximum A Posteriori denoiser with a regular associated prior. Again the result is stated for more general priors than TV prior. This proves that TV-LSE is a kind of regularization of TV-MAP denoiser, and would give another reason why staircasing effect cannot occur with a LSE risk [Nikolova 2004]. No closed formulation was found for the obtained smooth prior, but numerical experiments of Chapter 4 will prove that it is different from classical regularization of TV-MAP and has more interesting denoising properties.

Let p be a probability distribution defined on \mathbb{R}^n , satisfying the following assumption.

Assumption 4 p is absolutely continuous with respect to Lebesgue's measure and logarithmically concave (written shortly log-concave), i.e. there exists a convex and coercive mapping f from \mathbb{R}^n to $\mathbb{R}^+ \cup \{+\infty\}$ (called potential) such that p has density $\frac{1}{Z}e^{-f}$.

Notice that if p has a density $\frac{1}{Z}e^{-f}$, where Z is, as usual, a normalizing constant, and if p satisfies Assumption 4, then its potential is such that $\{y \mid f(y) < \infty\}$ has a non-empty interior (otherwise $\int_{\mathbb{R}^n} dp = 0$).

Let S_{LSE} denote the LSE-denoiser associated to the prior p . It is defined by

$$S_{LSE} : \begin{cases} \mathbb{R}^n & \rightarrow & \mathbb{R}^n \\ y & \mapsto & \frac{\int x \exp(-\frac{\|x-y\|^2}{2\sigma^2})p(x)dx}{\int \exp(-\frac{\|x-y\|^2}{2\sigma^2})p(x)dx}. \end{cases}$$

Remark We can assume $\sigma = 1$ by the following change of variables $\bar{y} = y/\sigma$, $\bar{x} = x/\sigma$, $\bar{p}(t) = p(\sigma t)$, $\bar{S}_{LSE} : y \mapsto S_{LSE}(\sigma y)/\sigma$, and rewrite

$$S_{LSE} : \begin{cases} \mathbb{R}^n & \rightarrow & \mathbb{R}^n \\ y & \mapsto & \frac{\int x \exp(-\frac{\|x-y\|^2}{2})p(x)dx}{\int \exp(-\frac{\|x-y\|^2}{2})p(x)dx}. \end{cases}$$

In the sequel σ is set to 1.

Theorem 3.3 *The following properties hold:*

1. S_{LSE} is continuously differentiable on \mathbb{R}^n . Its differential dS_{LSE} at every point y is a symmetric definite positive matrix.
2. S_{LSE} is monotone in the sense of Brezis [1966], i.e.

$$\forall y_1, y_2 \in \mathbb{R}^n, \langle S_{LSE}(y_2) - S_{LSE}(y_1), y_2 - y_1 \rangle \geq 0.$$

3. S_{LSE} is non-expansive, i.e.

$$\forall y_1, y_2 \in \mathbb{R}^n, \|S_{LSE}(y_2) - S_{LSE}(y_1)\| \leq \|y_2 - y_1\|$$

4. S_{LSE} is injective.

Proof : 1. The proof for $S_{LSE} \in \mathcal{C}^1$ is the same as in the TV-case. In the proof of Proposition 3.3, replace TV by f and everything works because f is assumed to be non-negative. The differential of S_{LSE} at point y satisfies, for each $h \in \mathbb{R}^n$

$$\begin{aligned} dS_{LSE}(y)(h) &= \frac{\int \langle h, x-y \rangle x e^{-\frac{\|x-y\|^2}{2}} p(x) dx}{\int \exp(-\frac{\|x-y\|^2}{2}) p(x) dx} - \frac{\int \langle h, x-y \rangle e^{-\frac{\|x-y\|^2}{2}} p(x) dx}{\int e^{-\frac{\|x-y\|^2}{2}} p(x) dx} \frac{\int x e^{-\frac{\|x-y\|^2}{2}} p(x) dx}{\int e^{-\frac{\|x-y\|^2}{2}} p(x) dx} \\ &= \frac{\int \langle h, x \rangle x e^{-\frac{\|x-y\|^2}{2}} p(x) dx}{\int \exp(-\frac{\|x-y\|^2}{2}) p(x) dx} - \frac{\int \langle h, x \rangle e^{-\frac{\|x-y\|^2}{2}} p(x) dx}{\int e^{-\frac{\|x-y\|^2}{2}} p(x) dx} \frac{\int x e^{-\frac{\|x-y\|^2}{2}} p(x) dx}{\int e^{-\frac{\|x-y\|^2}{2}} p(x) dx}. \end{aligned}$$

The differential $dS_{LSE}(y)$ can be interpreted as a covariance matrix

$$dS_{LSE}(y) = \mathbb{E}[\Xi_y {}^t \Xi_y] - \mathbb{E}\Xi_y \mathbb{E} {}^t \Xi_y = \text{Cov}\Xi_y \quad \text{with } \Xi_y \sim q_y,$$

where q_y has density $\frac{1}{Z} e^{-\frac{\|x-y\|^2}{2}} p(x)$. Indeed, for each $h \in \mathbb{R}^n$,

$$\begin{aligned} (\text{Cov}\Xi_y)h &= \mathbb{E}[\Xi_y {}^t \Xi_y h] - \mathbb{E}\Xi_y \mathbb{E} [{}^t \Xi_y h] \\ &= \mathbb{E}[\langle h, \Xi_y \rangle \Xi_y] - \mathbb{E}\langle h, \Xi_y \rangle \mathbb{E}\Xi_y, \end{aligned}$$

where we can recognize $dS_{LSE}(y)(h)$. In particular, $dS_{LSE}(y)$ is symmetric with non-negative eigenvalues. Let us prove now that $dS_{LSE}(y)$ is positive-definite. Let us assume that there exists a vector $h \neq 0$ in the kernel of $dS_{LSE}(y)$, i.e. such that

$$(\text{Cov}\Xi_y)h = 0.$$

Then multiplying on the left by ${}^t h$ yields

$${}^t h (\text{Cov}\Xi_y)h = \text{var}\langle h, \Xi_y \rangle = 0.$$

But the support of distribution q_y satisfies

$$\text{Supp}(q_y) = \text{Supp}(p) = \{y \in \mathbb{R}^n \mid f(y) < \infty\},$$

which has non-empty interior. Then $\langle h, \Xi_y \rangle$ cannot have a zero variance. We obtain a contradiction. Finally $dS_{LSE}(y)$ is a symmetric positive-definite matrix.

2. (adaptation of the proof of [Brezis 1966, Prop. 2]). Let y_1 and y_2 in \mathbb{R}^n , and let

$$\varphi(t) = S_{LSE}((1-t)y_1 + ty_2).$$

Then the derivative of φ satisfies

$$\varphi'(t) = dS_{LSE}((1-t)y_1 + ty_2)(y_2 - y_1).$$

As $dS_{LSE}((1-t)y_1 + ty_2)$ is a positive matrix, we get

$$\langle \varphi'(t), y_2 - y_1 \rangle = \langle dS_{LSE}((1-t)y_1 + ty_2)(y_2 - y_1), y_2 - y_1 \rangle \geq 0,$$

and the mapping $\psi(t) = \langle \varphi(t), y_2 - y_1 \rangle$ having derivative $\langle \varphi'(t), y_2 - y_1 \rangle$ is non-decreasing. In particular $\psi(1) \geq \psi(0)$, which writes

$$\langle S_{LSE}(y_2) - S_{LSE}(y_1), y_2 - y_1 \rangle \geq 0.$$

3. As S_{LSE} is differentiable, it suffices to prove that for any $y \in \mathbb{R}^n$, all the eigenvalues of $dS_{LSE}(y)$ are in the interval $[0, 1]$. We already proved (first part of the proposition) that the eigenvalues of $dS_{LSE}(y)$ are all positive. To complete the proof, we need the two following lemmas. Their proof will be given after the current proof.

Lemma 3.3 Let G denote the centered Gaussian kernel with standard deviation (s.d.) 1. Let $g = -\log(p * G)$. Then for all $y \in \mathbb{R}^n$,

$$dS_{LSE}(y) = I_n - H(g)(y). \quad (3.12)$$

where I_n is the identity matrix of \mathbb{R}^n , and $H(g)(y)$ is the Hessian matrix of g .

Lemma 3.4 $g = -\log(p * G)$ is convex.

Assume that Lemma 3.3 and Lemma 3.4 hold. As g is convex, $H(g)$ contains non-negative eigenvalues only. This proves that the eigenvalues of $dS_{LSE}(y)$ are less than 1.

4. Assume that $S_{LSE}(y_1) = S_{LSE}(y_2)$. Then considering again the mapping

$$\psi(t) = \langle S_{LSE}((1-t)y_1 + ty_2), y_2 - y_1 \rangle$$

satisfying $\psi(0) = \psi(1)$, its derivative

$$\psi'(t) = \langle dS_{LSE}((1-t)y_1 + ty_2)(y_2 - y_1), y_2 - y_1 \rangle$$

must vanish at a certain point $t_0 \in [0, 1]$. But $dS_{LSE}((1-t_0)y_1 + t_0y_2)$ is a positive-definite matrix, and consequently $\psi'(t) > 0$ provides a contradiction. \square

Proof of Lemma 3.3 : The LSE denoiser can be written using convolution products

$$S_{LSE}(y) = \frac{(Xp) * G}{p * G}, \quad (3.13)$$

where G denote the centered Gaussian kernel with s.d. 1, and $X : x \mapsto x$ is the monomial function with degree 1, i.e. the identity of \mathbb{R}^n (we use this notation because we shall need polynomials with higher degree in the sequel, it should not be confused with a random vector). For $i \in \{1, \dots, n\}$, X_i will denote the mapping $x \mapsto x_i$. As both G and $p * G$ are differentiable, and as $\frac{\partial G}{\partial x_i} = -X_i G$, the equality

$$p * \frac{\partial G}{\partial x_i} = \frac{\partial(p * G)}{\partial x_i}$$

can be derived into

$$\forall y \in \mathbb{R}^n, \quad (X_i p) * G(y) = y_i(p * G)(y) + \frac{\partial(p * G)}{\partial x_i}(y). \quad (3.14)$$

Now let $i \neq j$. As $\frac{\partial^2 G}{\partial x_i \partial x_j} = X_i X_j G$, the expression $\int (x_i - y_i)(x_j - y_j)p(x)G(y - x)dx$ can be derived into

$$\begin{aligned} & (X_i X_j p) * G(y) - y_i(X_j p) * G(y) - y_j(X_i p) * G(y) + y_i y_j (p * G)(y) \\ &= \int (x_i - y_i)(x_j - y_j)p(x)G(y - x)dx \\ &= \int p(x) \frac{\partial^2 G}{\partial x_i \partial x_j}(y - x)dx \\ &= \frac{\partial^2(p * G)}{\partial x_i \partial x_j}(y). \end{aligned} \quad (3.15)$$

Combining Equations (3.14) and (3.15) yields for $\Xi \sim q_y$

$$\begin{aligned} \text{cov}(\Xi_i, \Xi_j) &= \frac{(X_i X_j p) * G}{p * G}(y) - \frac{(X_i p) * G}{p * G}(y) \frac{(X_j p) * G}{p * G}(y) \\ &= \frac{\frac{\partial^2(p * G)}{\partial x_i \partial x_j}}{p * G}(y) - \frac{\frac{\partial(p * G)}{\partial x_i}}{p * G}(y) \frac{\frac{\partial(p * G)}{\partial x_j}}{p * G}(y) \\ &= -\frac{\partial^2 g}{\partial x_i \partial x_j}(y) \quad \text{where } g = -\log(p * G). \end{aligned}$$

If $i = j$, we must add another term because $\frac{\partial^2 G}{\partial x_i^2} = (X_i^2 - 1)G$, which implies

$$\begin{aligned} &(X_i^2 p) * G(y) - 2y_i(X_i p) * G(y) + y_i^2(p * G)(y) \\ &= \int (x_i - y_i)^2 p(x) G(y - x) dx \\ &= \int p(x) \left[\frac{\partial^2 G}{\partial x_i^2}(y - x) + G(y - x) \right] dx \\ &= (p * G)(y) + \frac{\partial^2(p * G)}{\partial x_i^2}(y), \end{aligned} \tag{3.16}$$

and then, combining Equations (3.14) and (3.16), for $\Xi \sim q_y$

$$\begin{aligned} \text{var}(\Xi_i) &= \frac{(X_i^2 p) * G}{p * G}(y) - \left(\frac{(X_i p) * G}{p * G} \right)^2 (y) \\ &= 1 + \frac{\frac{\partial^2(p * G)}{\partial x_i^2}}{p * G}(y) - \frac{\frac{\partial(p * G)}{\partial x_i}}{p * G}(y) \frac{\frac{\partial(p * G)}{\partial x_i}}{p * G}(y) \\ &= 1 - \frac{\partial^2 g}{\partial x_i^2}(y) \quad \text{where } g = -\log(p * G). \end{aligned}$$

We get $dS_{LSE}(y) = Id - H(g)(y)$, where

$$H(g)(y) = \left(\frac{\partial^2 g}{\partial x_i \partial x_j}(y) \right)_{ij}$$

is the Hessian matrix of g at point y . □

Proof of Lemma 3.4 : This is a direct consequence of [Prékopa 1973, Theorem 7], stating that the distribution given by the convolution product of two log-concave distributions is log-concave. □

Theorem 3.3 (items 2. and 3., i.e. monotony and non-expansiveness) are properties which could be required in the generalization of thresholding operators to high dimensional spaces. Moreover the theorem proves that S_{LSE} has a certain regularity and stability, and is injective: from a denoised image \hat{u}_{LSE} , it is theoretically possible to retrieve the noisy image v such that $\hat{u}_{LSE} = S_{LSE}(v)$.

A more restrictive condition on the prior allows S_{LSE} to be onto, as next proposition states.

Proposition 3.6 *If the potential $f = -\log p - \log Z$ of the prior is Lipschitz continuous, i.e.*

$$\exists k > 0, \quad \forall x, y \in \mathbb{R}^n, \quad |f(y) - f(x)| \leq k\|y - x\|,$$

then $S_{LSE} - I_n$ is bounded. Furthermore, S_{LSE} is onto, and S_{LSE} is a \mathcal{C}^1 -diffeomorphism of \mathbb{R}^n .

Proof : We first prove that $S_{LSE} - I_n$ is bounded. Let S^{n-1} denote the unit sphere of \mathbb{R}^n . Then let us prove that

$$\exists c > 0 \quad / \quad \forall u \in S^{n-1}, \forall y \in \mathbb{R}^n, \langle S_{LSE}(y) - y, u \rangle \leq c. \quad (3.17)$$

Let $u \in S^{n-1}$ and $y \in \mathbb{R}^n$. Then

$$\begin{aligned} \langle S_{LSE}(y) - y, u \rangle &= \frac{\int \langle x - y, u \rangle e^{-\frac{\|x-y\|^2}{2}} p(x) dx}{\int e^{-\frac{\|x-y\|^2}{2}} p(x) dx} \\ &= \frac{\int \langle x, u \rangle e^{-\frac{\|x\|^2}{2}} p(x+y) dx}{\int e^{-\frac{\|x\|^2}{2}} p(x+y) dx} \\ &= \frac{\int_{t \in \mathbb{R}} t e^{-\frac{t^2}{2}} \int_{\varepsilon \in u^\perp} e^{-\frac{\|\varepsilon\|^2}{2}} p(y+tu+\varepsilon) d\varepsilon dt}{\int_{t \in \mathbb{R}} e^{-\frac{t^2}{2}} \int_{\varepsilon \in u^\perp} e^{-\frac{\|\varepsilon\|^2}{2}} p(y+tu+\varepsilon) d\varepsilon dt} \quad \text{with } x = tu + \varepsilon \\ &\leq \frac{\int_{t \in \mathbb{R}} |t| e^{-\frac{t^2}{2}} \int_{\varepsilon \in u^\perp} e^{-\frac{\|\varepsilon\|^2}{2}} p(y+tu+\varepsilon) d\varepsilon dt}{\int_{t \in \mathbb{R}} e^{-\frac{t^2}{2}} \int_{\varepsilon \in u^\perp} e^{-\frac{\|\varepsilon\|^2}{2}} p(y+tu+\varepsilon) d\varepsilon dt} \end{aligned}$$

Now as $-\log p = f + \log Z$ is Lipschitz continuous, we can use the inequalities

$$p(y+tu)e^{-k\|\varepsilon\|} \leq p(y+tu+\varepsilon) \leq p(y+tu)e^{k\|\varepsilon\|}$$

to bound the inner integral $\int_{\varepsilon \in u^\perp} e^{-\frac{\|\varepsilon\|^2}{2}} p(y+tu+\varepsilon) d\varepsilon$ from above and below

$$p(y+tu) \int_{\varepsilon \in u^\perp} e^{-\frac{\|\varepsilon\|^2}{2} + 2k\|\varepsilon\|} d\varepsilon \leq \int_{\varepsilon \in u^\perp} e^{-\frac{\|\varepsilon\|^2}{2}} p(y+tu+\varepsilon) d\varepsilon \leq p(y+tu) \int_{\varepsilon \in u^\perp} e^{-\frac{\|\varepsilon\|^2}{2} - 2k\|\varepsilon\|} d\varepsilon,$$

which writes

$$b_k p(y+tu) \leq \int_{\varepsilon \in u^\perp} e^{-\frac{\|\varepsilon\|^2}{2}} p(y+tu+\varepsilon) d\varepsilon \leq a_k p(y+tu)$$

where a_k and b_k are positive numbers, independent from u and y . Then we have

$$\langle S_{LSE}(y) - y, u \rangle \leq \frac{a_k \int |t| e^{-\frac{t^2}{2}} p(y+tu) dt}{b_k \int e^{-\frac{t^2}{2}} p(y+tu) dt}.$$

Now, using again the fact that $-\log p$ is Lipschitz continuous, we have

$$p(y)e^{-kt} \leq p(y+tu) \leq p(y)e^{kt}$$

since u is a unit vector. Hence we get

$$\langle S_{LSE}(y) - y, u \rangle \leq \frac{a_k \int |t| e^{-\frac{t^2}{2}} p(y) e^{kt} dt}{b_k \int e^{-\frac{t^2}{2}} p(y) e^{-kt} dt} = \frac{a_k \int |t| e^{-\frac{t^2-2kt}{2}} dt}{b_k \int e^{-\frac{t^2+2kt}{2}} dt}$$

Letting $c = \frac{a_k \int |t| e^{-\frac{t^2-2kt}{2}} dt}{b_k \int e^{-\frac{t^2+2kt}{2}} dt} < \infty$ proves (3.17).

Now as $S_{LSE} - I_n$ is bounded, it is straightforward that

$$\lim_{\|y\| \rightarrow \infty} \frac{|\langle S_{LSE}(y), y \rangle|}{\|y\|} = +\infty. \quad (3.18)$$

If the dimension n is equal to 1, then S_{LSE} is continuous and $S_{LSE} - I_n$ is bounded, and thanks to the intermediate value theorem, S_{LSE} is onto. Now if the dimension n satisfies $n \leq 2$, we can use [Brezis 1966, Corollary 16], implying that if E is a vector space with finite dimension ≥ 2 , if S_{LSE} is a monotone and continuous mapping from E to E and if (3.18) holds, then S_{LSE} is onto. This concludes the first part of the proof.

To prove that S_{LSE} is a \mathcal{C}^1 -diffeomorphism, one only needs to check that S_{LSE}^{-1} is \mathcal{C}^1 . But Theorem 3.3 (1) claimed that for every $y \in \mathbb{R}^n$, $dS_{LSE}(y)$ is a positive-definite matrix. Then S_{LSE}^{-1} is continuous and differentiable with differential

$$dS_{LSE}^{-1}(y) = (dS_{LSE}(S_{LSE}^{-1}(y)))^{-1}$$

which is continuous because $y \mapsto dS_{LSE}(S_{LSE}^{-1}(y))$ is continuous. \square

Proposition 3.6 has consequences over TV-LSE denoising in practice. Indeed, TV-LSE denoising satisfies the assumptions of the proposition (see Section 3.4), and thus we expect the TV-LSE denoised images to contain more noise than the TV-MAP denoised images. Let us explain this. As S_{LSE} is a diffeomorphism, it should be possible to retrieve the noisy image from the denoised one alone, and this implies that some information about the noise is kept in the denoised image one way or another. Most probably this information lies in residual noise in the denoised image. Actually numerical experiments held in Chapter 4 generally show some residual noise. This gives to the denoised images a much more natural aspect than the ones denoised by TV-MAP.

Now that some general properties concerning LSE denoisers have been settled, a comparison to MAP denoisers

$$S_{MAP} : \begin{cases} \mathbb{R}^n & \rightarrow \mathbb{R}^n \\ y & \mapsto \arg \min_x \frac{\|x-y\|^2}{2} + f(x) \end{cases}$$

called *proximity operators* in the literature (see [Moreau 1962; Combettes and Wajs 2004, 2005]) can be held. S_{MAP} can be proven to be non-expansive [Moreau 1965; Combettes and Wajs 2004] and monotone, but is not always continuously differentiable nor injective, as the TV case proves. Function f can be viewed as the potential associated to a certain prior distribution p_{MAP} .

We shall prove that there exists a log-concave prior p_{MAP} , such that the MAP-denoiser associated to p_{MAP} is identically equal to the LSE-denoiser S_{LSE} associated to p . The next theorem expounds this fact with potentials instead of priors. We first need an additional assumption.

Assumption 5 *The set $S_{LSE}(\mathbb{R}^n)$ is convex.*

Thanks to Proposition 3.6, Assumption 5 is satisfied by the TV prior, because in this case $S_{LSE}(\mathbb{R}^n) = \mathbb{R}^n$ is convex. Assumption 5 is probably satisfied for any log-concave prior distribution p .

Theorem 3.4 *Let f satisfy Assumptions 4 and 5. Then there exists f_{MAP} convex on $S_{LSE}(\mathbb{R}^n)$ and \mathcal{C}^2 , such that*

$$\forall y \in \mathbb{R}^n, S_{LSE}(y) = \arg \min_x \frac{\|x - y\|^2}{2} + f_{MAP}(x).$$

If an explicit formulation of potential f_{MAP} was reachable, faster algorithms than the one developed in Chapter 2, such as forward-backward algorithms [Combettes and Wajs 2005], would be possible. Besides this may give clues for the design of new image priors. Now we only have a very rough approximation of it (see the very end of the section). Therefore the real impact of the theorem lies in the regularity of f_{MAP} : even if the initial potential f is not regular, the MAP potential f_{MAP} will always be regular. Therefore the main “morality” of the theorem is that the notion of prior is vague, and that the choice of the MAP or LSE risk can be crucial in a large class of estimation problems. Besides, more specifically, this result gives another clue for the absence of staircasing in TV-LSE denoised images. Indeed, a result from [Nikolova 2004] which is the converse of Theorem 1.4 cited in Chapter 1 Subsection 1.2.3 states that if f_{MAP} is \mathcal{C}^∞ , then (1.4) can almost never occur. Our framework only proves f_{MAP} to be \mathcal{C}^2 , which prevents us to directly apply this result. However this work can be viewed as a first attempt to generalize the result stating the absence of staircasing, to more general prior distributions than p .

Before proving Theorem 3.4, let us consider the following lemma.

Lemma 3.5 *Let $C = S_{LSE}(\mathbb{R}^n)$. The mapping $S_{LSE}^{-1} - I_n$ defined on C can be identified to an exact differential form.*

Proof of Lemma 3.5: Let $y \in C$. Then $(S_{LSE}^{-1} - I_n)(y)$ is in \mathbb{R}^n , and denoting ω the differential form

$$\omega(y) = ((S_{LSE}^{-1} - I_n)(y))_1 dx_1 + \dots + ((S_{LSE}^{-1} - I_n)(y))_n dx_n$$

(where x_i denotes the i -th component of a vector x), $(S_{LSE}^{-1} - I_n)$ can be identified to the differential form ω . Furthermore, thanks to Assumption 5, C is convex hence simply connected. Now, thanks to Poincaré’s theorem, this differential form is exact iff it is closed, i.e. iff

$$\forall i \neq j, \quad \frac{\partial(S_{LSE}^{-1} - I_n)(y)_i}{\partial y_j} = \frac{\partial(S_{LSE}^{-1} - I_n)(y)_j}{\partial y_i},$$

i.e. iff the differential $dS_{LSE}^{-1}(y)$ is a symmetric matrix. But Theorem 3.3 already stated that $dS_{LSE}(y)$ was a symmetric positive definite matrix. Thus the inverse matrix is also symmetric. But as $dS_{LSE}^{-1}(S_{LSE}(x)) = dS_{LSE}(x)^{-1}$, choosing $x = S_{LSE}^{-1}(y)$, we get the closedness of $(S_{LSE}^{-1} - I_n)$. \square

Proof of Theorem 3.4: Thanks to Lemma 3.5, $S_{LSE}^{-1} - I_n$ can be interpreted as the gradient of a certain function. f_{MAP} will be that function, whose construction is as follows:

- Outside C , set $f_{MAP}(y) = +\infty$.
- Fix $y_0 \in C$, and set $f_{MAP}(y_0) = 0$.
- For every $y \in C$, let γ be a curve linking y_0 to y , i.e. a \mathcal{C}^1 mapping

$$\gamma : [0, 1] \rightarrow C \quad \text{such that } \gamma(0) = y_0 \text{ and } \gamma(1) = y.$$

Then $f_{MAP}(y)$ is defined by

$$f_{MAP}(y) = f_{MAP}(y_0) + \int_0^1 \langle (S_{LSE}^{-1} - I_n)(\gamma(t)), \gamma'(t) \rangle dt$$

and does not depend on the choice of γ .

Now let us check that f_{MAP} is such that

$$\forall y \in \mathbb{R}^n, S_{LSE}(y) = \arg \min_x \frac{\|x - y\|^2}{2} + f_{MAP}(x).$$

By construction we have

$$\forall y \in C, \quad \nabla f_{MAP}(y) = S_{LSE}^{-1}(y) - y. \quad (3.19)$$

Therefore the minimizer x of $\frac{\|x-y\|^2}{2} + f_{MAP}(x)$ makes the gradient of this energy equal to 0, and hence satisfies

$$x - y + S_{LSE}^{-1}(x) - x = 0,$$

i.e. $x = S_{LSE}(y)$.

Now let us check that f_{MAP} is \mathcal{C}^2 . S_{LSE} is a \mathcal{C}^1 -diffeomorphism from \mathbb{R}^n to C . Then $\nabla f_{MAP}(y) = S_{LSE}^{-1}(y) - y$ is \mathcal{C}^1 . Hence f_{MAP} is \mathcal{C}^2 .

Now let us check that f_{MAP} is convex. Let us show that the Hessian matrix of f_{MAP} is definite positive at all points. Let $y \in C$. Thanks to Equation (3.19), the Hessian matrix of f_{MAP} satisfies

$$\forall i, j, \quad (H(f_{MAP}))_{i,j}(y) = \frac{\partial}{\partial x_i} (S_{LSE}^{-1} - I_n)_j(y) = \frac{\partial S_{LSE}^{-1,j}}{\partial x_i}(y) - y_j.$$

But denoting $x = S_{LSE}^{-1}(y)$, we have

$$\left(\frac{\partial S_{LSE}^{-1}}{\partial x_i}(y) \right)_{1 \leq i, j \leq n} = dS_{LSE}^{-1}(y) = (dS_{LSE}(x))^{-1},$$

and it follows that $dS_{LSE}(x)$ is a symmetric positive-definite matrix with eigenvalues in $(0, 1]$ (see Theorem 3.3). Then $(dS_{LSE}(x))^{-1}$ is symmetric positive-definite with eigenvalues in $[1, +\infty)$. We get

$$H(f_{MAP})(y) = (dS_{LSE}(x))^{-1} - I_n$$

where $(dS_{LSE}(x))^{-1} - I_n$ is a symmetric non-negative matrix. Hence $H(f_{MAP})(y)$ is non-negative for every $y \in C$. That is why f_{MAP} is convex on C .

By the way, Equations (3.17) and (3.19) put together yield the fact that if f is Lipschitz continuous, then ∇f_{MAP} is bounded. It means that f_{MAP} is also Lipschitz continuous. \square

f_{MAP} is now proven to be convex, but also needs to be coercive, i.e. to satisfy

$$f_{MAP}(x) \xrightarrow{\|x\| \rightarrow \infty} +\infty,$$

to be regarded as the potential of a probability distribution (i.e. for $e^{-f_{MAP}}$ to be integrable). The proof for coercivity requires another assumption.

Assumption 6 $p * G$ is strictly log-concave.

Remark about Assumption 6 A result from [Prékopa 1973] states that the convolution of two log-concave probability density functions is log-concave. This result is deduced from Brunn-Minkowski inequality [Prékopa 1973; Maurey 2003]. However we found no paper in the literature about log-concave measures concerning strict log-concavity. In our case, a result like “the convolution of a log-concave p.d.f. defined on \mathbb{R}^n with a strictly log-concave p.d.f. (a Gaussian function) is strictly log-concave” would be enough for Assumption 6 to always be true. This result holds in the case where both density functions are continuous, positive and defined on \mathbb{R} (courtesy of Gilles Pagès). A generalization to higher dimensions could be the object of future investigation.

Proposition 3.7 Under Assumption 6, f_{MAP} is coercive.

The proof for Proposition 3.7 begins with a fixed point lemma.

Lemma 3.6 S_{LSE} has a fixed point, i.e.

$$\exists x_0 \in \mathbb{R}^n, \quad S_{LSE}(x_0) = x_0.$$

Proof of Lemma 3.6 : Equations (3.13) and (3.14) put together imply that S_{LSE} can be written as a gradient

$$\forall y \in \mathbb{R}^n, \quad S_{LSE}(y) = \nabla(I_n - g)(y) = y - \nabla g(y),$$

with $g = -\log(p * G)$ being differentiable. The existence of a fixed point for S_{LSE} is equivalent to the existence of $x_0 \in \mathbb{R}^n$ such that

$$\nabla g(x_0) = 0.$$

As g is convex (Lemma 3.4), such a x_0 corresponds to a minimum for g . Now, in order to prove that this minimum is reached in \mathbb{R}^n , let us prove that g is coercive, i.e. that

$$p * G(y) \xrightarrow{\|y\| \rightarrow \infty} 0.$$

This result is very natural since $p * G$ is both L^1 and very regular, but we detail the proof. Let $\varepsilon > 0$ be a positive real number. As p is integrable and non-negative, there exists $R > 0$ such that

$$\int_{B(0,R)^c} p(t) dt \leq \varepsilon,$$

where $B(0, R)$ denotes the ball centered at 0 with radius R , and $B(0, R)^c$ its complement. Then for all $y \in \mathbb{R}^n$, the convolution $p * G(y)$ can be decomposed into

$$p * G(y) = \int_{B(0,R)} p(t)G(y-t) dt + \int_{B(0,R)^c} p(t)G(y-t) dt. \quad (3.20)$$

The first term goes to 0 when $\|y\| \rightarrow \infty$. Indeed,

$$\int_{B(0,R)} p(t)G(y-t) dt \leq \int_{B(0,R)} p(t) dt \max_{B(y,R)} G \leq \max_{B(y,R)} G,$$

and the Gaussian function G satisfies $\max_{B(y,R)} G \rightarrow 0$ when $\|y\| \rightarrow \infty$. The second term of (3.20) satisfies

$$\int_{B(0,R)^c} p(t)G(y-t) dt \leq \int_{B(0,R)^c} p(t) dt \max_{B(y,R)^c} G \leq \varepsilon \cdot G(0).$$

Now putting things in a nutshell, let $M > 0$ be such that

$$\|y\| > M \Rightarrow \max_{B(y,R)} G \leq \varepsilon.$$

Then for any y with $\|y\| > M$, $p * G(y)$ is bounded by

$$p * G(y) \leq \max_{B(y,R)} G + \varepsilon \cdot G(0) \leq (1 + G(0)) \varepsilon,$$

which proves the convergence of $p * G(y)$ towards 0 when $\|y\| \rightarrow \infty$, i.e. the coercivity of g .

As g is convex, coercive and differentiable, its minimum is reached at some point $x_0 \in \mathbb{R}^n$, which is a fixed point for the operator S_{LSE} . \square

Proof of Proposition 3.7 : Assumption 6 strengthens Lemma 3.4 by assuming that the mapping g defined by $g = -\log(p * G)$ is strictly convex and not only convex. Let us recall from Lemma 3.3 that

$$\forall y \in \mathbb{R}^n, \quad dS_{LSE}(y) = I_n - H(g)(y).$$

As g is \mathcal{C}^2 and strictly convex, the spectrum of $H(g)(y)$ is inside $(0, +\infty)$. Now as S_{LSE} is monotone and one-to-one, the spectrum of $dS_{LSE}(y)$ is inside $(0, 1)$, and the spectrum of $dS_{LSE}^{-1}(y)$ is inside $(1, +\infty)$ for any $y \in \mathbb{R}^n$. These facts imply

$$\forall y \in \mathbb{R}^n, \quad \text{Sp}(d(S_{LSE}^{-1} - I_n)(y)) \subset (0, +\infty).$$

Let x_0 be a fixed point for S_{LSE} (Lemma 3.6), and u a unit vector. Then for each $T > 0$, we have

$$\begin{aligned} \langle \nabla f_{MAP}(x_0 + Tu), u \rangle &= \langle (S_{LSE}^{-1} - I_n)(x_0 + Tu), u \rangle \\ &= \langle (S_{LSE}^{-1} - I_n)(x_0 + Tu) - (S_{LSE}^{-1} - I_n)(x_0), u \rangle \\ &= \int_0^T \langle d(S_{LSE}^{-1} - I_n)(x_0 + tu)(u), u \rangle dt, \end{aligned}$$

where $d(S_{LSE}^{-1} - I_n)(x_0 + tu)$ is a definite positive matrix, and therefore $\langle \nabla f_{MAP}(x_0 + Tu), u \rangle$ is > 0 , which implies the coercivity of f_{MAP} . \square

Under Assumption 6, Proposition 3.7 holds, and $e^{-f_{MAP}}$ is integrable. Therefore there exists a distribution p_{MAP} whose potential is f_{MAP} , and such that S_{MAP} associated to p_{MAP} equals S_{LSE} associated to $p = p_{LSE}$.

Let us go back to the case of an arbitrary noise variance σ^2 . The previous results can be adapted, and it is mentionable that when σ is small, the distribution p_{MAP} satisfies

$$\frac{\nabla p_{MAP}}{p_{MAP}} = \frac{\nabla(p_{LSE} * G_\sigma)}{p_{LSE} * G_\sigma} + O(\sigma^2). \quad (3.21)$$

Indeed, Equation (3.14) yields

$$S_{LSE} = I_n + \sigma^2 \frac{\nabla(p_{LSE} * G_\sigma)}{p_{LSE} * G_\sigma}.$$

In the meantime, as f_{MAP} is \mathcal{C}^2 , p_{MAP} is differentiable, and S_{MAP} satisfies

$$(I_n - \sigma^2 \frac{\nabla p_{MAP}}{p_{MAP}}) \circ S_{MAP} = I_n$$

which is the Euler equation attached to the minimization of $x \mapsto \|x - y\|^2 + 2\sigma^2 f_{MAP}(x)$. For σ small enough, the function $(I_n - \sigma^2 \frac{\nabla p_{MAP}}{p_{MAP}})$ can be inverted thanks to the global inversion theorem, and S_{MAP} writes

$$S_{MAP} = (I_n - \sigma^2 \frac{\nabla p_{MAP}}{p_{MAP}})^{-1}.$$

Making S_{LSE} and S_{MAP} equal yields

$$(I_n - \sigma^2 \frac{\nabla p_{MAP}}{p_{MAP}}) \circ (I_n + \sigma^2 \frac{\nabla(p_{LSE} * G_\sigma)}{p_{LSE} * G_\sigma}) = I_n$$

and then

$$\frac{\nabla(p_{LSE} * G_\sigma)}{p_{LSE} * G_\sigma} - \frac{\nabla p_{MAP}}{p_{MAP}} = \sigma^2 \frac{\nabla p_{MAP}}{p_{MAP}} \circ \frac{\nabla(p_{LSE} * G_\sigma)}{p_{LSE} * G_\sigma}.$$

In particular, when σ^2 goes to 0, we have (3.21), and p_{MAP} can be roughly approximated by

$$p_{MAP} \approx p_{LSE} * G_\sigma.$$

Chapter 4

Numerical experiments

Résumé Ce chapitre regroupe des images débruitées par la méthode TV-LSE, calculées par l'algorithme décrit dans le chapitre 2. Nous discutons le critère de convergence dans sa pratique, et nous montrons que la convergence visuelle pour une image de taille 512×512 est atteinte en un temps raisonnable. L'influence des deux paramètres du modèle est étudiée, et s'avère compatible avec les résultats asymptotiques prouvés au chapitre 3. Le débruitage TV-LSE est aussi comparé à d'autres méthodes de débruitage, parmi lesquels la minimisation TV (TV-MAP) et les moyennes non-locales. En plus de l'absence effective du phénomène de *staircasing*, les images débruitées par TV-LSE combinent les qualités d'un débruitage homogène et une bonne préservation de la géométrie, ce qui leur confère un rendu très naturel. D'autres expériences sont menées sur des signaux synthétiques et sur des images naturellement bruitées.

Abstract In this chapter, images denoised by TV-LSE are shown, where TV-LSE has been implemented as described in Chapter 2. The practical convergence criterion is discussed, and it is shown that visual convergence can be achieved for 512×512 images within reasonable computation time. The influence of the 2 parameters of TV-LSE is shown, and is consistent with the asymptotic results of Chapter 3. TV-LSE is also compared to several other denoising methods, among which TV-minimization and Non-Local means. The resulting images combine homogeneous noise reduction and geometry preservation, still avoiding staircasing effect and leading to very natural rendering. Experiments are also shown on synthetic signals and on images degraded by natural noise.

In this chapter, numerical experiments about TV-LSE denoising are presented. The algorithms presented in Chapter 2 have been implemented in C under Megawave2 (free software for image processing ¹) on a laptop with a processor at 2.0 GHz. We give some details here about the implementation.

- The numerical scheme for *TV* was chosen as the one of [Chambolle 2004] allowing a fast algorithm for TV-MAP, since the comparison between TV-LSE and TV-MAP is compulsory.

¹<http://megawave.cmla.ens-cachan.fr/>

If the domain Ω is the rectangle $\{1, \dots, n_1\} \times \{1, \dots, n_2\}$, then for any pixel $(x, y) \in \Omega$, the two components of the gradient of an image u is defined by

$$\nabla u^1(x, y) = \begin{cases} u(x+1, y) - u(x, y) & \text{if } x < n_1 \\ 0 & \text{if } x = n_1 \end{cases}$$

$$\nabla u^2(x, y) = \begin{cases} u(x, y+1) - u(x, y) & \text{if } y < n_2 \\ 0 & \text{if } y = n_2. \end{cases}$$

This scheme for the gradient offers a good compromise between isotropy and stability. The total variation of an image u is then defined by

$$TV(u) = \frac{1}{n_1 n_2} \sum_{x=1}^{n_1} \sum_{y=1}^{n_2} \sqrt{|\nabla u^1(x, y)|^2 + |\nabla u^2(x, y)|^2}.$$

This standard scheme for total variation is a special case of $TV_{\varepsilon, p}$ defined in Chapter 3 Section 3.4 with $\varepsilon = 0$, $\mathcal{N}_{(x, y)} = \{(x, y), (x+1, y), (x, y+1)\}$ and the norms

$$\begin{cases} \|\cdot\|_k & = L^2\text{-norm on } \mathcal{N}_k \text{ if } \mathcal{N}_k \subset \Omega \\ \|\cdot\|_k & = L^1\text{-norm on } \mathcal{N}_k \text{ if } \mathcal{N}_k \cap \Omega^c \neq \emptyset. \end{cases}$$

Hence it satisfies all the assumptions needed for both the algorithm convergence (results of Chapter 2) and the theoretical results of Chapter 3, i.e. essentially non-negativity, convexity, finiteness, shift invariance and Lipschitz continuity.

- The MCMC algorithm is run after a scaling tuning step (Algorithm 4). First Algorithm 4 is run with 2 chains initialized by a random uniform noise on $[0, 255]$. The last state of both chains is recovered and used as the initialization of (U_0) and (\tilde{U}_0) for Algorithm 3. Although Algorithm 4 already contains a burn-in step, in practice the burn-in optimization procedure in Algorithm 3 is essential to allow better convergence rate, because the optimal burn-in parameter depends on the total number of iterations and rises above the initial burn-in, in practice.

4.1 Discussion on the stopping criterion

At first sight numerical experiments on TV-LSE denoising seem prohibitive in terms of computation time. We will try here to show that for a large range of parameters (λ, σ) a relatively small number of iterations is enough to make the algorithm converge. A 512×512 image can then be denoised by TV-LSE within a few minutes or even a few seconds sometimes.

We shall use $|\Omega|$ -subsamped MCMCs and the stopping criterion described in Chapter 2 Section 2.3. One iteration will then refer to $|\Omega|$ small iterations of the complete chain (Υ_n) . After choosing a positive parameter ε , we run the two MCMCs U_n and \tilde{U}_n and stop when their respective averages satisfy

$$\|S_n^b - \tilde{S}_n^b\| \leq 2\varepsilon$$

for which we expect the estimate $\frac{S_n^b + \tilde{S}_n^b}{2}$ of \hat{u}_{LSE} to satisfy

$$\left\| \frac{S_n^b + \tilde{S}_n^b}{2} - \hat{u}_{LSE} \right\| \leq \varepsilon.$$

ε should be chosen as little as possible, but taking it too small yields high computation time because of the larger number of iterations then required. As Figure 4.1 shows, the choice $\varepsilon = 1$ for images with gray values in $[0, 255]$ seems sufficient for a good visual quality while maintaining reasonable computation time. Indeed, smaller ε yield denoised images which visually are absolutely similar (zoomed details are shown on Figure 4.2). The precision $\varepsilon = 1$ is reached here within 8 minutes for this 512×512 image, while much more time would be necessary to reach fractions of 1. A parallelization in the programming seems possible: first the two chains could be computed separately and compared every R iterations; secondly a local processing of the image could yield a task splitting into many local tasks. Hence one could manage to have much shorter computation times with a graphic card using GPU (Graphic Processor Unit) techniques for instance [Owens et al. 2007; Borghi et al. 2008].

The graph of Figure 4.3 represents $\|S_n^b - \tilde{S}_n^b\|/2$ (i.e. ε if we stopped at iteration n) as a function of n . It converges to zero when $n \rightarrow \infty$ as $C \cdot n^{-0.66}$ approximately. The rate of convergence is then quite slow, but small n can be enough to reach reasonable values of ε .

Of course the number of iterations needed to reach the stopping criterion depends on the parameters λ and σ . Briefly speaking it is linked to the spread of the posterior distribution $\frac{1}{Z}e^{-E_\lambda/(2\sigma^2)}$. If the posterior distribution has a large spread (large σ , small λ) then the exploration of the posterior distribution will take a long time and a large number of iterations will be necessary to reach the wanted ε . Conversely when the posterior distribution is very concentrated around its maximum (small σ , large λ) then the exploration can be done within few iterations thanks to an efficient burn-in. But we must bear in mind the fact that our stopping criterion is only valid for negligible bias situations. When the posterior distribution is very concentrated, the bias can be preponderant over the variance in the bias-variance decomposition of the risk (see Equation (2.11)), and even an efficient burn-in procedure cannot always prevent the algorithm from stopping before convergence.

In practice we never encountered convergence problems with the algorithm of Chapter 2 as soon as $\sigma \geq \lambda/10$. As $\sigma \approx 0$ corresponds to the case where TV-LSE behaves like TV-MAP (i.e. TV-minimization of [Rudin et al. 1992]), a good initialization in the algorithm allows a correct convergence control even in the case $\sigma \leq \lambda/10$.



Figure 4.1: Influence of ε in the stopping criterion (I). The Lena image (with gray levels inside $[0, 255]$) noised by a Gaussian white noise with standard deviation (s.d.) $\sigma = 10$ (top-left image) is denoised using TV-LSE with the algorithm explained in Chapter 2. The three other images are the partial averages $(S_n^b + \tilde{S}_n^b)/2$ of a unique run of the algorithm, for three values of n . The parameters are all fixed ($\sigma = 10$, $W = 30$, $\alpha = 27$, no burn-in) except from the number of iterations n . The top-right image, named $\hat{u}_{LSE}(300)$, obtained by a 300-iteration run, is such that the distance between the 2 chains of the algorithm $\|S_{300}^0 - \tilde{S}_{300}^0\|$ is closest to 2, so that we expect the image to be at a distance $\varepsilon \approx \|S_{300}^0 - \tilde{S}_{300}^0\|/2 \approx 1$ from the true TV-LSE denoised image \hat{u}_{LSE} (actually $n = 300$ corresponds to the multiple of 50 reaching the distance closest to 2, because the images of the MCMCs were saved only every 50 images). Waiting for some more iterations yields the bottom-left image $\hat{u}_{LSE}(800)$ ($\varepsilon \approx 0.5$ with $n = 800$), and the bottom-right image $\hat{u}_{LSE}(5000)$ ($\varepsilon \approx 0.2$ with $n = 5000$). The three images look the same, it is difficult to make them apart, even on the zooms presented on next figure. For this 512×512 image, one must count 166 seconds, i.e. almost 3 minutes to run 100 iterations of the algorithm (for the two 512^2 -subsampled chains).




Figure 4.2: Influence of ε in the stopping criterion (II). Details of the noisy Lena image (first column), and TV-LSE denoised images with $n = 300$ ($\varepsilon \approx 1$, second column), $n = 800$ ($\varepsilon \approx 0.5$, third column), and $n = 5000$ ($\varepsilon \approx 0.2$, last column). The contrast has been enhanced in order to make the comparison easier. On the first row is shown the 50×50 window which maximizes the L^2 -distance between $\hat{u}_{LSE}(300)$ cropped and $\hat{u}_{LSE}(5000)$ cropped. The second row shows the 50×50 window which maximizes the L^2 -distance between $\hat{u}_{LSE}(800)$ cropped and $\hat{u}_{LSE}(5000)$ cropped. The third row shows a window containing contrasted edges, the last row shows another window containing texture. In all cases, eye inspection reveals almost no difference. Consequently keeping $\varepsilon \approx 1$ is visually enough for our denoising tasks, and allows reasonable computation times (on the whole 512×512 image 8 minutes instead of 2h20min for $\varepsilon = 0.2$).




Figure 4.3: Influence of ε in the stopping criterion (III). Half of the L^2 -distance between the 2 MCMCs of TV-LSE (i.e. ε) is plotted against the number n of iterations. The curve behaves like $C \cdot n^{-0.66}$. $n = 300$ results in a convergence with precision $\varepsilon \approx 1$: visually the image will not be different from the one found with a much higher n .

4.2 Influence of parameters σ , λ and β

Contrary to TV-MAP (TV-minimization of [Rudin et al. 1992]) which depends on parameter λ only, our algorithm deals with 2 parameters λ and σ . σ^2 stands for the noise variance, and λ is selected in such a way that $\frac{1}{Z}e^{-\lambda TV/(2\sigma^2)}$ is a correct prior for images.

In most experimental purposes the noise variance is known, or at least can be estimated [Rousseuw and Leroy 1987; Kervrann and Boulanger 2006]. It seems reasonable then to choose this known or estimated value for σ^2 . However as Theorem 3.1 states, TV-MAP corresponds to the limit case where $\sigma^2 = 0$, and this scheme does denoise images, even if the noise is statistically modeled as zero. Furthermore real denoising applications do not necessarily require a complete noise removal but only need to sweep out part of the noise, to make sure that the information contained in the image is not distorted. Hence it does not look absurd to consider σ as a free parameter of the model.

Consequently λ and σ are considered as two regularizing parameters of the denoiser, where σ controls the amount of noise and λ balances the regularity term versus the data-fidelity term in the energy E_λ . Instead of considering (λ, σ) one could equivalently consider (β, σ) as the set of parameter of the denoiser, where

$$\beta = \frac{\lambda}{2\sigma^2}$$

is the inverse of the temperature in the Gibbs field defined by the prior

$$\pi(u) = \frac{1}{Z}e^{-\beta TV(u)}.$$

β is probably a better regularity parameter than λ , and is more coherent with the constrained version of TV-LSE (see Chapter 1 Subsection 1.3.1), but taking (β, σ) for parameter set makes comparison between TV-LSE and TV-MAP harder for a given shape of energy. This is why we will maintain both parameter sets (λ, σ) or (β, σ) in our experiments.

In the following figures (Figures 4.4 to 4.13) the aim is to get an intuitive comprehension of the parameters λ , σ and β .

- Figure 4.4 shows TV-LSE denoised images for different parameters λ and σ . The regularity of the denoised image increases when λ increases or σ decreases. Therefore both λ and σ appear as regularizing parameters. But λ acts more like a TV-regularizing parameter (when λ increases, texture is removed, but edges remain sharp) whereas σ controls the distance to the noisy image v and to the MAP-denoised image $\hat{u}_{MAP}(\lambda)$, because thanks to Theorem 3.1,

$$\begin{cases} \hat{u}_{LSE}(\lambda, \sigma) \xrightarrow{\sigma \rightarrow 0} \hat{u}_{MAP}(\lambda) \\ \hat{u}_{LSE}(\lambda, \sigma) \xrightarrow{\sigma \rightarrow +\infty} v. \end{cases}$$

- Figure 4.5 also shows TV-LSE denoised images for different parameters λ and σ , applied on a noisy image containing thin stripes (a detail of Barbara image). It illustrates the fact that TV-LSE denoising manages to remove noise while maintaining the oscillating pattern, yet reducing the contrast. Even if staircasing theoretically cannot occur in TV-LSE denoised images (see Section 3.3), some staircasing effect is visible yet for small enough σ , because

human eyes can interpret close gray level pixels as one constant region. Taking a larger σ gives visually staircasing-free images.

- The other set of parameters (β, σ) is used in Figure 4.6, where TV-LSE denoised images are shown for different values of the parameters β and σ . The image contains more noise when σ and β are small, and its regularity increases when one of the parameters increases. β acts as a TV-regularizing parameter when σ is small (first small and contrasted details such as cars are preserved and the the image becomes practically constant), and as a blurring parameter when σ is large (the contrasted details shrink, and the edges are smoothed).



Figure 4.4: Influence of the parameters λ and σ on $\hat{u}_{LSE}(\lambda, \sigma)$. The Lena image was corrupted by a Gaussian white noise with s.d. 10 (see Figure 4.1, top-left image), and denoised using the fully automatic TV-LSE algorithm (optimal scaling, automatic burn-in, and $\varepsilon = 1$). A column considered alone (σ fixed) starts on top with a rather noisy image, and is more regular as we get down (as λ increases), with preserved boundaries, regularized smooth regions (no staircasing visible) and erased texture. A row considered alone (λ fixed) starts on the left with a regularized image which is not far from $\hat{u}_{MAP}(\lambda)$ (because $\hat{u}_{LSE}(\lambda, \sigma)$ converges to $\hat{u}_{MAP}(\lambda)$ when $\sigma \rightarrow 0$), and details of the texture and then noise appear when scrolling to the right, hence coming closer to the noisy image (the input). These observations are consistent with the table of Figure 3.2. SNR are given Figure 4.7.




Figure 4.5: Influence of (λ, σ) on $\hat{u}_{LSE}(\lambda, \sigma)$. The same experiment as in Figure 4.4 is presented on Barbara. The first column, corresponding to $\sigma = 2$, is very close to TV-MAP denoised images (because $\hat{u}_{LSE}(\lambda, \sigma)$ converges to $\hat{u}_{MAP}(\lambda)$ when $\sigma \rightarrow 0$), and the stripes get brutally erased, beginning with the thinner stripes, when scrolling down the column, giving space to constant regions. Conversely, for the other columns, the stripes are progressively erased, with intermediary states where the stripes' contrast is reduced. The artificial-looking constant regions are avoided when taking σ large enough. SNR are given Figure 4.7.




Figure 4.6: Influence of (β, σ) on \hat{u}_{LSE} . A detail of the Nîmes image from the CNES was noised by a 10-s.d. Gaussian white noise (the obtained image is visually equivalent to the top-left image), and denoised using the fully automatic algorithm (optimal scaling, automatic burn-in, and $\varepsilon = 1$). A column considered alone (σ fixed) starts on top (β small) with a noisy image, and the image gets more regular as β increases. But this regularity is more TV-regularity when σ is small (i.e. sharp edges separate constant regions and contrasted details such as cars are preserved), whereas it is more a blur-regularity when σ is large (i.e. the contrasted details shrink, the edges become blurry, and the image gets smooth). Contrast has been linearly enhanced to make details more visible. SNR are given Figure 4.7.

	$\sigma = 0$	$\sigma = 2$	$\sigma = 10$	$\sigma = 50$
$\lambda = 2$	14.93	14.84	14.49	13.82
$\lambda = 8$	18.50	18.17	17.01	14.64
$\lambda = 32$	17.57	16.46	18.52	17.11
$\lambda = 128$	12.47	9.30	10.27	16.57

images of Fig. 4.4

	$\sigma = 0$	$\sigma = 2$	$\sigma = 10$	$\sigma = 50$
$\lambda = 2$	14.57	14.46	14.45	14.31
$\lambda = 8$	13.71	13.78	13.93	14.52
$\lambda = 32$	6.85	7.00	7.70	11.68
$\lambda = 128$	4.62	4.46	4.56	6.28

images of Fig. 4.5

	$\sigma = 0$	$\sigma = 2$	$\sigma = 10$	$\sigma = 50$
$\beta = 0.01$	9.38	9.40	10.23	12.17
$\beta = 0.1$	9.38	9.82	13.08	3.91
$\beta = 1$	9.38	13.24	3.24	0.20
$\beta = 10$	9.38	3.97	0.25	0.005

images of Fig. 4.6

Figure 4.7: Tables of SNR for images of Figures 4.4, 4.5 and 4.6, with $\text{SNR} = 10 \log_{10} \frac{\sigma_u^2}{\|\hat{u}_{LSE} - u\|^2}$, where σ_u^2 is the empirical variance of the original image u . The columns corresponding to $\sigma = 0$ were computed via Chambolle's algorithm for TV-MAP for the 2 first tables, and via the convergence $\hat{u}_{LSE}(\lambda, \sigma) \rightarrow v$ when $(\lambda, \sigma) \rightarrow 0$ concerning the third table.

Now it seems useful to consider TV-LSE denoising when all the parameters are fixed but one, and to appreciate the evolution of the denoised image when the last parameter changes.

- First of all on Figure 4.8 we consider σ fixed, and λ free. It equivalent to considering σ fixed and β free. According to the first table of Figure 3.2, when σ is fixed

$$\begin{cases} \hat{u}_{LSE} \xrightarrow{\lambda \rightarrow 0} v \\ \hat{u}_{LSE} \xrightarrow{\lambda \rightarrow \infty} \bar{v}\mathbf{1}, \end{cases}$$

where v is the noisy image, \bar{v} its gray level average, and $\mathbf{1}$ the constant image equal to 1 on each pixel. Between these extreme attitudes $\hat{u}_{LSE}(\lambda, \sigma)$ goes through images where noise is reduced, edges are preserved and smooth regions are preserved (no staircasing).

- Now on Figure 4.9 we consider the case where λ , i.e. the ratio between the data-fidelity term and the regularity term in the energy E_λ , is fixed, and σ takes all positive values. Again according to the table of Figure 3.2, we have, when λ is fixed,

$$\begin{cases} \hat{u}_{LSE} \xrightarrow{\sigma \rightarrow 0} \hat{u}_{MAP}(\lambda) \\ \hat{u}_{LSE} \xrightarrow{\sigma \rightarrow \infty} v, \end{cases}$$

and indeed for small values of σ the denoised image contains practically no noise, the edges are well contrasted, the texture is somewhat erased, and staircasing is visible. Therefore when σ increases, edges remain well contrasted, and the image remains noise-free, while staircasing disappears: the images look more natural than for small values of σ . Finally for large values of σ noise appears together with some thin texture: the images are close to the noisy input v .

- Figure 4.10 deals with the case where $\beta = \frac{\lambda}{2\sigma^2}$ is fixed and σ is free. Theory (Figure 3.2) forecasts the following convergences

$$\begin{cases} \hat{u}_{LSE}(\beta, \sigma) \xrightarrow{\sigma \rightarrow 0} v \\ \hat{u}_{LSE}(\beta, \sigma) \xrightarrow{\sigma \rightarrow \infty} \bar{v}\mathbf{1}. \end{cases}$$

Indeed in practice a very small σ yields an image close to the noisy input v . When σ increases the image becomes more regular: noise is reduced, edges are preserved, then small details in the texture disappear. When σ is large, the texture is erased, the edges become blurry; the image is then smooth and melts down into a constant image.

- The previous situations all make the image go from a noisy image to a regularized image. The nature of regularization is particular to each case, but the comparison of the denoised images is a biased comparison between more or less noisy images. A more interesting case is when the level of denoising, or *method noise*

$$\|\hat{u}_{LSE}(\lambda, \sigma) - v\|,$$

is fixed, and a parameter is free. The mapping $\lambda \mapsto \|\hat{u}_{LSE}(\lambda, \sigma) - v\|$ is increasing for a fixed σ with values in $(0, \|v - \bar{v}\mathbf{1}\|)$. This is why when $\|\hat{u}_{LSE}(\lambda, \sigma) - v\|$ and σ are fixed with $\|\hat{u}_{LSE}(\lambda, \sigma) - v\|$ inside the interval $(0, \|v - \bar{v}\mathbf{1}\|)$, there exists exactly one value $\lambda(\sigma)$ such that $\hat{u}_{LSE}(\lambda(\sigma), \sigma)$ reaches the desired level of denoising. In Figure 4.11 we chose a level of denoising in $(0, \|v - \bar{v}\mathbf{1}\|)$, and for several values of σ , we displayed the TV-LSE denoised images $\hat{u}_{LSE}(\lambda(\sigma), \sigma)$, where $\lambda(\sigma)$ attains the required denoising level. When σ is small the image is close to a TV-MAP denoised image, and contains many staircasing artefacts. When σ is large the image is a bit blurry and still contains noise. Then for a given level of denoising, the parameter σ tunes an equilibrium between staircasing artefacts reduction and blur control.

The table in Figure 4.12 gives the SNR of the images displayed on Figures 4.8 to 4.11. It confirms the fact that medium values of σ , λ , β achieve a better image reconstruction, and that TV-LSE is worth being compared to other denoising methods.




Figure 4.8: $\sigma = 10$ is fixed, and TV-LSE denoising is run for different values of λ ($\lambda \in \{2, 4, 8, 16, 32, 64, 128\}$, and $\sigma = 10$ corresponds to the s.d. of the noise). It is equivalent to fixing σ and letting $\beta = \frac{\lambda}{2\sigma^2}$ take values in \mathbb{R}^+ . When λ (or β) is small, the denoised image \hat{u}_{LSE} starts by being very close to the noisy image v . Then when λ (or β) increases the noise disappears, the homogeneous regions are smoothed out with no staircasing. Then the texture is erased, and the last image is not far from being a piecewise smooth image.

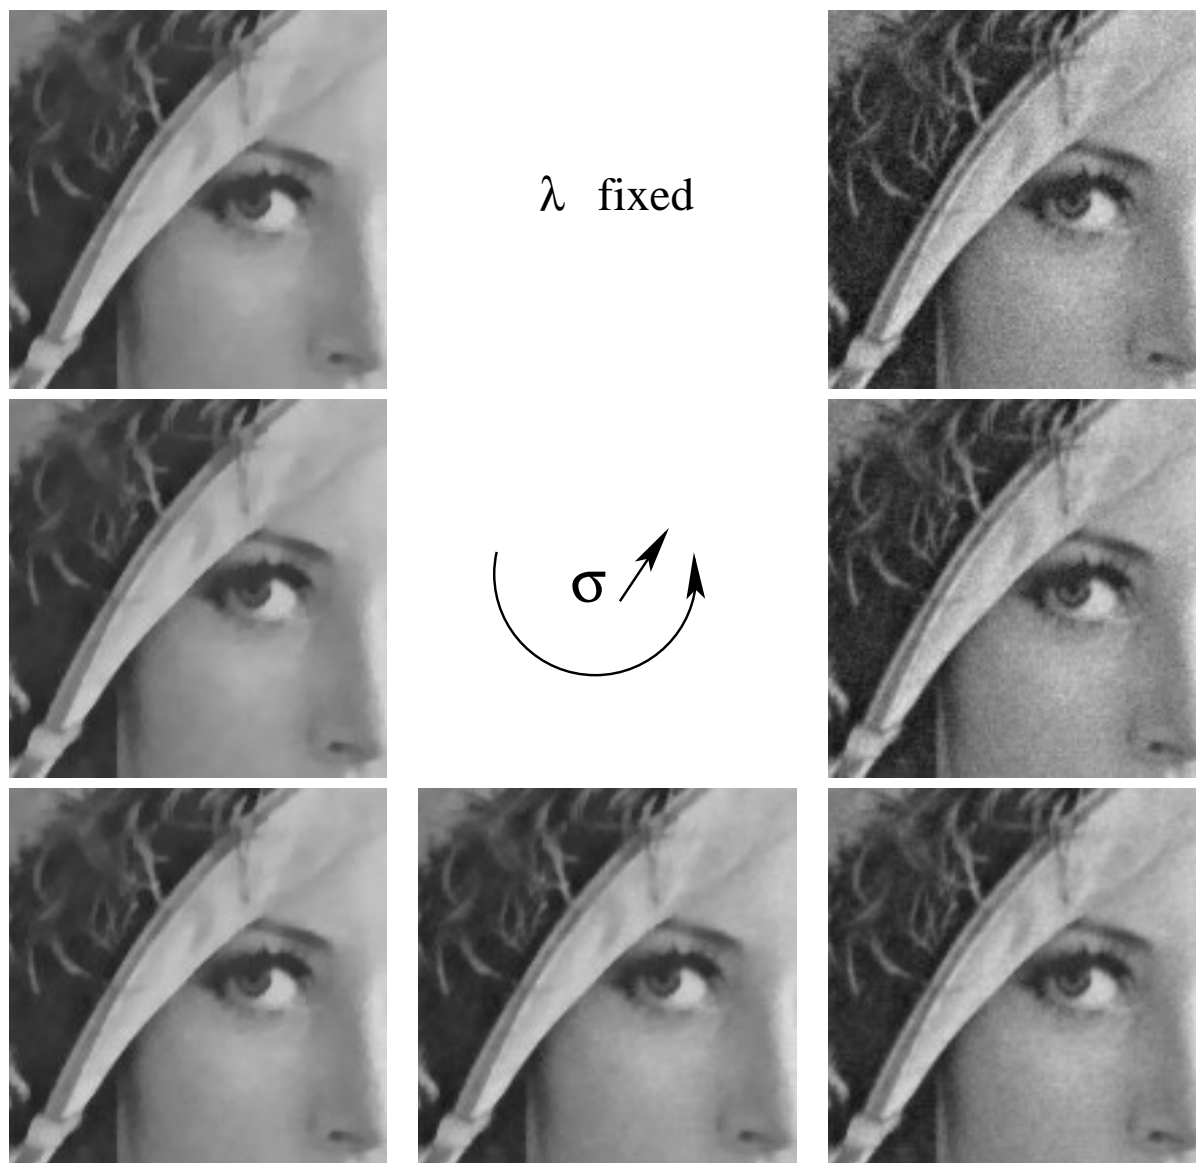


Figure 4.9: $\lambda = 30$ is fixed, and TV-LSE denoising is run for different values of σ ($\sigma \in \{1.25, 2.5, 5, 10, 20, 40, 80\}$ and the noise has s.d. 10). When σ is small, the denoised image \hat{u}_{LSE} is practically noise-free, but some texture is erased and some staircasing is visible: Lena's cheek and hat contain boundaries which do not exist in the original Lena image. This image is very close to the MAP-denoised image $\hat{u}_{MAP}(\lambda)$, to which $\hat{u}_{LSE}(\lambda, \sigma)$ converges when $\sigma \rightarrow 0$. Then when σ increases, staircasing effect disappears, texture comes back together with noise, and \hat{u}_{LSE} looks more and more like the noisy image. This is consistent with the convergence $\hat{u}_{LSE}(\lambda, \sigma) \rightarrow v$ when $\sigma \rightarrow \infty$.




Figure 4.10: $\beta = \frac{\lambda}{2\sigma^2} = 0.15$ is fixed, and TV-LSE denoising is run for different values of σ ($\sigma \in \{1.25, 2.5, 5, 10, 20, 40, 80\}$ and the noise has s.d. 10). For small values of σ , the denoised image is close to the noisy image. When σ increases the image is regularized, the edges are preserved while texture is erased, and then the denoised image blurs and becomes completely smooth. σ then acts as a blurring parameter when β is fixed, for large enough values of σ .




Figure 4.11: The level of denoising $\|\hat{u}_{LSE}(\lambda, \sigma) - v\| = 9$ is fixed, and TV-LSE denoising is run for different values of σ ($\sigma \in \{0, 1.25, 2.5, 5, 10, 20, 40\}$ and the noise has s.d. 10). For a given σ there exists at most one value of λ or β such that the level of denoising is reached. When σ increases, when they exist, the λ corresponding to the level of denoising increases too, while β decreases. $\sigma = 0$ gives rise to a TV-MAP denoised image: the image is free from noise, but some texture is erased, and staircasing is flagrant on the cheek for instance. When σ increases, staircasing disappears and the denoised image becomes blurry. Then for a given level of denoising, the parameters σ , λ and β tune a compromise between blurring and staircasing.

	$\lambda = 2$	$\lambda = 4$	$\lambda = 8$	$\lambda = 16$	$\lambda = 32$	$\lambda = 64$	$\lambda = 128$
SNR	14.50	15.41	17.02	18.87	18.53	14.97	10.32
$\ \hat{u}_{LSE} - u\ $	9.02	8.12	6.75	5.45	5.67	8.54	14.59
$\ \hat{u}_{LSE} - v\ $	1.63	2.63	4.44	6.96	9.55	12.48	17.52

images of Fig. 4.8 ($\sigma = 10$)

	$\sigma = 1.25$	$\sigma = 2.5$	$\sigma = 5$	$\sigma = 10$	$\sigma = 20$	$\sigma = 40$	$\sigma = 80$
SNR	16.72	16.88	17.55	18.70	18.67	17.48	15.94
$\ \hat{u}_{LSE} - u\ $	6.98	6.86	6.34	5.56	5.58	6.40	7.64
$\ \hat{u}_{LSE} - v\ $	11.01	10.93	10.48	9.33	7.58	5.34	3.35

images of Fig. 4.9 ($\lambda = 30$)

	$\sigma = 1.25$	$\sigma = 2.5$	$\sigma = 5$	$\sigma = 10$	$\sigma = 20$	$\sigma = 40$	$\sigma = 80$
SNR	13.85	14.75	17.65	18.70	14.18	9.60	6.20
$\ \hat{u}_{LSE} - u\ $	9.71	8.76	6.27	5.56	9.36	15.85	23.44
$\ \hat{u}_{LSE} - v\ $	0.98	1.96	5.20	9.33	13.17	18.66	25.49

images of Fig. 4.10 ($\beta = 0.15$)

	$\sigma = 0.625$	$\sigma = 1.25$	$\sigma = 2.5$	$\sigma = 5$	$\sigma = 10$	$\sigma = 20$	$\sigma = 40$
SNR	19.38	19.59	19.49	19.41	18.94	18.38	17.77
$\ \hat{u}_{LSE} - u\ $	5.13	5.02	5.07	5.12	5.41	5.77	6.18
$\ \hat{u}_{LSE} - v\ $	9.06	9.00	9.08	8.94	9.05	9.01	9.07

images of Fig. 4.11 ($\|v - \hat{u}_{LSE}\|_2 \approx 9$)

Figure 4.12: Tables with the SNR and the denoising level of the denoised images whose details were shown in Figures 4.8 to 4.11. The 3 first tables correspond to images going from the noisy image to an over-regularized version of the image; it seems natural that the optimal SNR be obtained for medium values of parameters. The last table corresponds to the case where the level of denoising is fixed, the parameter σ takes different positive values, and λ is tuned to get the desired level of denoising with precision 1% (see the last row). Here again the optimal SNR is reached by a medium-valued σ , achieving a trade-off between staircasing and blur.

Before coming to comparisons between TV-LSE and other classical denoising techniques, Figure 4.13 gives an illustration about the location of $\hat{u}_{LSE}(\lambda, \sigma)$ relatively to the noisy image v and the TV-MAP denoised images $\hat{u}_{MAP}(\lambda)$. For each value of λ , the images $(\hat{u}_{LSE}(\lambda, \sigma))_{\sigma>0}$ run inside an arch put up between v and $\hat{u}_{MAP}(\lambda)$. Several conjectures can be deduced from this figure. First, the curves never cross the x -axis, which shows that TV-LSE denoised images are not aligned with the original image and their corresponding TV-MAP denoised images: $\hat{u}_{LSE}(\lambda, \sigma)$ is not a simple deviation from $\hat{u}_{MAP}(\lambda)$ in the direction of the noisy input v . Furthermore $\hat{u}_{LSE}(\lambda, \sigma)$ seems to live inside the ball having diameter $[v, \hat{u}_{MAP}(\lambda)]$, property which happened to be too hard to prove. Last it can be noticed that the ends of the curves badly approximate v or $\hat{u}_{MAP}(\lambda)$: the deviation is constantly orthogonal to the x -axis. This is most probably due to the stopping criterion of our algorithm and to the high dimensionality of the framework.




Figure 4.13: This graph represents a projection in a 2-D space of the images $(\hat{u}_{LSE}(\lambda, \sigma))$. The original noisy image v was placed at the origin of the graph. On the x -axis, some TV-MAP denoised images $\hat{u}_{MAP}(\lambda)$ were placed according to their distance to v (small discs). The thin curves are each associated to a particular value of λ , and represent the denoised images $(\hat{u}_{LSE}(\lambda, \sigma))$ for several values of $\sigma > 0$, according to their distance d_1 to the original noisy image v , and their distance d_2 to $\hat{u}_{MAP}(\lambda)$. When σ is small $\hat{u}_{LSE}(\lambda, \sigma)$ is close to $\hat{u}_{MAP}(\lambda)$, while when σ is large $\hat{u}_{LSE}(\lambda, \sigma)$ is close to v . The convergences are not perfect on the graph because the precision ε of the estimation of \hat{u}_{LSE} was taken positive in the algorithm. The errors of the limit images (e.g. extremal points of the $(\lambda = 4)$ -curve) seem to be orthogonal to the $(v, \hat{u}_{MAP}(\lambda))$ -axis, which can be explained by the high dimension of the framework: almost every vector taken in a ball centered at v or $\hat{u}_{MAP}(\lambda)$ is orthogonal to $(v, \hat{u}_{MAP}(\lambda))$. The bold curve is an arc of circle centered at v with radius 9, representing the set of images satisfying $\|\hat{u}_{LSE} - v\| = 9$, i.e. corresponding to a level of denoising equal to 9. For λ large enough, there always exists some σ such that $\hat{u}_{LSE}(\lambda, \sigma)$ attains some given level of denoising.

4.3 Denoising normalization and comparisons to other methods

The TV-LSE denoiser has to be compared to classical denoising methods. Every method depends on one or several parameters tuning the level of denoising. In order to obtain quantitative comparability of images we must tune the parameters of the methods to achieve equal levels of denoising. For a TV-LSE denoised image $\hat{u}_{LSE}(\lambda, \sigma)$ and some *level of denoising*, or *method noise*

$$d = \|\hat{u}_{LSE}(\lambda, \sigma) - v\|,$$

we shall keep for comparison the methods below, with the tuning of parameters explained:

- **TV-MAP** denoising (our main benchmark), i.e. Total Variation minimization of Rudin et al. [1992], from which TV-LSE has been constructed. The algorithm used for the minimization of the energy

$$E_{\lambda_{MAP}}(u) = \|u - v\|^2 + \lambda_{MAP} TV(u)$$

is Chambolle's [Chambolle 2004]. The parameter λ_{MAP} controls regularity of the image, and is tuned in such a way that the denoised image achieves the required level of denoising:

$$\|\hat{u}_{MAP}(\lambda_{MAP}) - v\| = d.$$

We expect to have $\lambda_{MAP} < \lambda$, because when taking $\lambda_{MAP} = \lambda$, $\hat{u}_{LSE}(\lambda, \sigma)$ seems to be inside the ball centered at v with radius $\|\hat{u}_{MAP}(\lambda) - v\|$ (see Figure 4.13), then the distance $\|\hat{u}_{MAP}(\lambda) - v\|$ seems to be always larger than $\|\hat{u}_{LSE}(\lambda, \sigma) - v\|$.

- **TV-barycentre** To be able to compare $\hat{u}_{LSE}(\lambda, \sigma)$ and $\hat{u}_{MAP}(\lambda)$ with the same λ (i.e. the same balance between regularity and data-fidelity term in E_λ), $\hat{u}_{MAP}(\lambda)$ will be linearly combined with the noisy image v via

$$\hat{u}_{bary} = t \hat{u}_{MAP}(\lambda) + (1 - t) v \quad \text{with } t = \frac{\|\hat{u}_{MAP}(\lambda) - v\|}{d}$$

and we obtain a barycentre of $\hat{u}_{MAP}(\lambda)$ and v which is at distance d from v . This choice is also motivated by the observation that the quality of denoising often increases both visually and in SNR when deviating the TV-MAP estimate towards v . Visual quality is better when the noise is not completely removed.

- **TV- ϵ** Before the paper of Chambolle [2004] which solved the exact TV-minimization problem iteratively but fastly, people used gradient descents of the energy E_λ , which required the introduction of an artificial smoothing parameter ϵ into the gradient norm's scheme, defined by

$$|\nabla u|_\epsilon^2(x) = \frac{\epsilon^2 + (a - b)^2 + (b - c)^2 + (c - d)^2 + (d - a)^2 + \frac{(a+b-c-d)^2 + (a-b-c+d)^2}{2}}{(2 + \sqrt{2})^2}$$

for instance. Taking $\epsilon > 0$ makes the energy

$$\|u - v\|^2 + \lambda_\epsilon \sum_x |\nabla u|_\epsilon(x)$$

differentiable and the gradient descent feasible, and avoids staircasing [Nikolova 2004]. But there are many drawbacks to this method: taking ϵ too small makes the algorithm run very slowly, and for any $\epsilon > 0$ the minimized energy is not the true total variation energy and loses its geometric interpretation (coarea formula and Cheeger sets [Alter et al. 2005a]).

The TV-LSE denoiser avoids staircasing while maintaining an exact formulation of total variation. Moreover in Section 3.5, LSE denoisers are proven to be MAP denoisers for a smoother prior distribution. We would like to check that this prior is not a simple ϵ -regularized TV prior. This is why \hat{u}_{LSE} is compared to the TV- ϵ denoised image \hat{u}_ϵ . ϵ will be frozen equal to 1 while the λ_ϵ will be tuned in order to have

$$\lambda_\epsilon \quad \text{such that} \quad \|\hat{u}_\epsilon - v\| = d.$$

- **Wiener filter** In the same vein as TV- ϵ , we consider the Wiener filter (in its simplest version) which consists in minimizing the energy

$$\|u - v\|^2 + \lambda_{Wiener} \sum_x |\nabla u|^2(x). \quad (4.1)$$

The Total variation has been replaced by the Sobolev H^1 -norm which gives infinite energy to discontinuous images: the filter tends to blur images. The underlying prior distribution $\frac{1}{Z}e^{-\beta \sum_x |\nabla u|^2(x)}$ is even smoother than the prior of TV- ϵ . The Wiener filter is linear and translation invariant, so it can be computed by a Fast Fourier Transform. Again the parameter λ_{Wiener} will be tuned in such a way that the minimizer \hat{u}_{Wiener} of (4.1) satisfies

$$\lambda_{Wiener} \quad \text{such that} \quad \|\hat{u}_{Wiener} - v\| = d.$$

- **MCM** The Euler equation associated to the minimization of E_λ is the PDE

$$\frac{\partial u}{\partial t} = \text{curv}(u) \quad \text{where} \quad \text{curv}(u) = \text{div} \frac{\nabla u}{|\nabla u|}$$

where the variable λ is changed into the stopping time T of the PDE. Changing this PDE into

$$\frac{\partial u}{\partial t} = |\nabla u| \text{curv}(u)$$

results in the well-known Mean Curvature Motion, which has the property of being multiscale analysis invariant, morphological invariant, translation invariant and rotation invariant [Alvarez et al. 1993; Guichard and Morel 2001]. Even if this PDE was originally built as a multiscale analysis tool, it can be used as an image denoiser. The stopping time T is linked to the scale of study, i.e. to the required quantity of simplification of the image. Let \hat{u}_{MCM} denote the denoised image. Once again the stopping time will be chosen as a

$$T \quad \text{such that} \quad \|\hat{u}_{MCM} - v\| = d.$$

The numerical method used here is taken from Megawave2. It is a consistent scheme achieving isometric and scale invariance for a variety of transforms leaving the grid invariant, and experimentally stable.

- **Gaussian convolution** The oldest denoising method is to make an image evolve with the heat equation, whose solution is a Gaussian convolution. A fast implementation can be done using Fast Fourier Transform, and the Gaussian kernel bandwidth is

$$g \quad \text{such that} \quad \|\hat{u}_{Gauss} - v\| = d.$$

- **Non-Local means (NL-means)** In the general goal of this thesis where we aim at linking Non-Local means (NL-means) [Buades et al. 2005] and TV-minimization, comparison of TV-LSE to NL-means seems compulsory. The link between the two is that the TV-LSE denoiser in its construction (Chapter 1 Subsection 1.3.2) can be interpreted as a diffusion as well as NL-means (Buades et al. [2007]; Lafon et al. [2006]; Peyré [2008]).

In the NL-means, each pixel x of an image v is denoised by estimating the true gray level $u(x)$ by an average of the gray values of pixels y belonging to a *large* window centered at x (hence the word Non-Local). For a given patch shape, this average is weighted by the similarity of the patch centered at x and the patch centered at each pixel y . More specifically, let \mathcal{N}_y denote a generic square patch centered at pixel y (here we shall take a 5×5 square as advised by Buades et al. [2005]), and \mathcal{W}_x a large square window centered at pixel x (here we shall consider a 21×21 window). Then the denoised image \hat{u}_{NL} is given by its gray levels

$$\hat{u}_{NL}(x) = \frac{\sum_{y \in \mathcal{W}_x} v(y) e^{-\frac{\|v(\mathcal{N}_y) - v(\mathcal{N}_x)\|_2^2}{2h^2}}}{\sum_{y \in \mathcal{W}_x} e^{-\frac{\|v(\mathcal{N}_y) - v(\mathcal{N}_x)\|_2^2}{2h^2}}},$$

where the distance between patches is chosen to be a classical normalized Euclidean distance

$$\|v(\mathcal{N}_y) - v(\mathcal{N}_x)\|_2^2 = \frac{1}{|\mathcal{N}_0|} \sum_{k \in \mathcal{N}_0} (v(y+k) - v(x+k))^2.$$

(\mathcal{N}_0 denotes an artificial patch centered at the origin). At this point the NL-means now only depends on parameter h , acting as the bandwidth of the Gaussian kernel on the distance between patches. h controls the level of denoising, since taking h large allows many patches $v(\mathcal{N}_y)$ to participate to the average added to the central patch $v(\mathcal{N}_x)$, which makes \hat{u}_{NL} move away from v .

The complete NL-means actually depends on 4 parameters: the bandwidth h , the patch size, the window size, and a parameter a in the distance between patches which can be smoothed by a factor $e^{-|k|^2/(2a^2)}$ to be less dependent on the borders of the patch. We are conscious that letting only one parameter free instead of 4 does not behave in favor of NL-means in our experiments. However h remains the most crucial parameter of NL-means and allows a precise control of the denoising level.

Despite its popularity, denoising by wavelet coefficient thresholding will not be taken into account in our experiments, because the method is considered too far from TV-LSE.




Figure 4.14: (I) Small level of denoising ($\|\hat{u} - v\| = 5.94$ for the Lena image noised by a Gaussian noise with s.d. = 10). Only part of the noise can be removed. First column: details of the noisy image. An affine contrast change (unique for each line) has been applied on all the images of Figures 4.14 to 4.16 in order to improve the visibility of details. This is why the gray levels are saturated on the noisy images. Second column: TV-LSE denoising with parameters $(\lambda, \sigma) = (20, 20)$. Noise remains but the global aspect is natural. Third column: TV-MAP denoising with parameter $\lambda_{MAP} = 7.46$. The remaining noise is mainly concentrated on isolated pixels. Staircasing is barely visible.




Figure 4.15: (II) Small level of denoising ($\|\hat{u} - v\| = 5.94$) (sequel of Figure 4.14). First column: TV-barycentre denoising with parameters $(\lambda, t) = (20, 0.61)$. The noise is reduced while a natural aspect is preserved. This shows that in this case TV-MAP is worth being noised again by the original noise. Second column: TV- ϵ denoising with $(\epsilon, \lambda_\epsilon) = (1, 8.28)$. Despite a positive ϵ the images are similar to the TV-MAP denoised images: contrasted isolated pixels with a starting staircasing effect. Third column: denoising by Wiener filter with $\lambda_{Wiener} = 0.74$. Noise is correctly reduced but the edges are blurred.




Figure 4.16: (III) Small level of denoising ($\|\hat{u} - v\| = 5.94$) (sequel of Figure 4.15). First column: denoising by Mean Curvature Motion, with parameter $T = 0.45$. Second column: denoising by Gaussian convolution with $g = 0.56$. As in Wiener denoising, noise is correctly reduced but the edges are blurred. Third column: NL-means, with parameter $h = 5.23$. The noise reduction is extremely heterogeneous. This results in unnatural looking images, despite excellent values of SNR (Figure 4.23).



Figure 4.17: (I) Middle level of denoising ($\|\hat{u} - v\| = 9.33$ for the Lena image noised by a Gaussian noise with s.d. = 10). An affine contrast change (unique for each line) has been applied on all the images of Figures 4.17 to 4.19 in order to improve the visibility of details. This is why some gray levels are saturated. First column: details of the noisy image. Second column: TV-LSE denoising with parameters $(\lambda, \sigma) = (30, 10)$. Texture is quite well recovered, as well as edges and homogeneous regions which really look natural. Some noisy texture remains. Third column: TV-MAP denoising with parameter $\lambda_{MAP} = 17.3$. A few contrasted isolated pixels are still visible, whereas staircasing becomes significant.




Figure 4.18: (II) Middle level of denoising ($\|\hat{u} - v\| = 5.94$) (sequel of Figure 4.17). First column: TV-barycentre denoising with $(\lambda, t) = (30, 0.87)$. The correction with the noisy image is not enough to conceal the staircasing effect due to MAP estimation. Besides the images still contain much noise. Second column: TV- ϵ with parameters $(\epsilon, \lambda_\epsilon) = (1, 19.32)$. The contrasted isolated pixels are attenuated, but the regularization by ϵ does not imply a drastic reduction of staircasing effect. Third column: denoising by Wiener filter with $\lambda_{Wiener} = 2.65$. Blur increases while a particular noisy texture covers all the homogeneous regions.




Figure 4.19: (III) Middle level of denoising ($\|\hat{u} - v\| = 5.94$) (sequel of Figure 4.18). First column: denoising by Mean Curvature Motion, with $T = 1.15$. The edges are well recovered but an embarrassing turbulent texture covers all the homogeneous regions. Second column: denoising by Gaussian convolution with $g = 0.83$. Noise is removed but all the details have vanished because of the blur. Third column: denoising by NL-means with $h = 8.47$. Edges and details are very well recovered. Still one can criticize the lack of homogeneity in the denoising which gives some artificial aspect to the detail images.



Figure 4.20: (I) High level of denoising ($\|\hat{u} - v\| = 13.06$ for the Lena image noised by a Gaussian noise with s.d. = 10). An affine contrast change (unique for each line) has been applied on all the images of Figures 4.20 to 4.22 in order to improve the visibility of details. First column: details of the noisy image. Second column: TV-LSE denoising with parameters $(\lambda, \sigma) = (100, 20)$. The noise has been removed. A very slight blur has erased some fine textures, but the global aspect is nice. Third column: TV-MAP denoising with parameter $\lambda_{MAP} = 73.1$. The noise has been removed while the edges remain well localized. However the ubiquitous staircasing effect over-simplifies the images and maintain an artificial aspect to these (the bottom image for instance is definitely damaged).



Figure 4.21: (II) High level of denoising ($\|\hat{u} - v\| = 13.06$) (sequel of Figure 4.17). First column: TV-barycentre denoising with $(\lambda, t) = (100, 0.93)$. As in Figure 4.18 the addition of noise does not suffice to hide the staircasing effect. However the global aspect is better than for TV-MAP: more details are visible and the global quality is somewhat better. Second column: TV- ϵ with parameters $(\epsilon, \lambda_\epsilon) = (1, 76.5)$. Compared to TV-MAP denoised images, staircasing is greatly reduced. Some fine details are erased. Third column: denoising by Wiener filter with $\lambda_{Wiener} = 11.160714$. Edges are blurry while noise persists.




Figure 4.22: (III) High level of denoising ($\|\hat{u} - v\| = 13.06$) (sequel of Figure 4.18). First column: denoising by Mean Curvature Motion, with $T = 6.8$. Noise is removed, but geometry of the images is highly distorted. Second column: denoising by Gaussian convolution with $g = 1.75$. The blur has erased all the fine structures. Third column: denoising by NL-means with $h = 25.2$. Quite a perfect denoising for the required level of denoising, but some details (eyelashes, large stripes of the hat) have vanished.

First the experiments are conducted on a single image (Lena) denoised with 3 different denoising levels (Figures 4.14 to 4.22). For a small level of denoising, noise is expected to be partially removed while maintaining the details of the image, while for a large level of denoising, noise is expected to be fully removed, and the precision of the remaining image is observed.

For a low level of denoising (Figures 4.14 to 4.16), TV-LSE achieves a good compromise between noise removal, preservation of the geometry of the image and homogeneity of denoising. Indeed, other TV-based denoising methods have spurious contrasted isolated pixels and staircasing starts being visible. NL-means denoising is terribly heterogeneous. For a medium level of denoising (Figures 4.17 to 4.19), TV-LSE greatly reduces the noise while no staircasing (unlike TV-MAP, TV-bary and TV- ϵ) nor blur (unlike Wiener and Gaussian convolution) is visible. TV-LSE does not create artificial texture, contrary to MCM and even NL-means. For a high level of denoising (Figures 4.20 to 4.22), the TV-LSE denoised image is a bit blurry but much lesser than Gaussian convolution and Wiener, and still avoids staircasing effect (unlike TV-MAP, TV-bary and TV- ϵ) and creation of spurious texture (unlike MCM). In all the denoised images, the TV-LSE denoised images have the particularity to look really natural, compared to most others.

For completeness a table of SNR is given in Figure 4.23. TV-LSE do not outperform the other methods in terms of SNR; anyway the SNR values are clearly decorrelated from visual quality (in these experiments for instance), and supervised comparisons of the zoomed denoised images are developed rather than extensive SNR tables.

	TV-LSE	TV-MAP	TV-bary	TV- ϵ	Wiener	MCM	Gauss.	NL-means
global	18.77	19.39	18.36	19.35	17.80	18.37	14.42	19.91
1st window	16.41	16.82	15.81	16.80	15.82	15.48	12.28	17.03
2nd window	14.54	15.16	14.00	15.07	13.64	14.57	9.52	15.48
3rd window	21.20	21.75	21.07	22.05	17.47	19.41	13.78	22.44
4th window	9.31	9.08	8.58	9.12	9.36	8.43	8.21	8.38

small level of denoising (Fig. 4.14 to 4.16)

	TV-LSE	TV-MAP	TV-bary	TV- ϵ	Wiener	MCM	Gauss.	NL-means
global	18.18	18.26	18.30	18.45	17.88	17.33	17.93	17.81
1st window	16.00	15.87	15.63	16.03	15.87	15.04	15.93	14.64
2nd window	15.31	15.29	15.19	15.47	15.17	14.81	15.25	14.04
3rd window	17.40	17.87	17.97	18.17	16.78	16.28	16.83	18.82
4th window	7.54	7.59	7.79	7.83	7.18	6.35	7.18	8.04

middle level of denoising (Fig. 4.17 to 4.19)

	TV-LSE	TV-MAP	TV-bary	TV- ϵ	Wiener	MCM	Gauss.	NL-means
global	15.08	14.42	13.88	14.47	14.27	14.68	14.45	14.82
1st window	12.50	12.01	11.14	11.83	11.82	10.24	12.32	12.10
2nd window	10.58	9.78	9.51	10.03	9.39	10.08	9.55	10.29
3rd window	18.52	18.97	18.10	19.52	13.23	20.52	13.82	21.58
4th window	8.10	6.75	6.95	6.78	8.56	7.37	8.23	6.41

large level of denoising (Fig. 4.20 to 4.22)

Figure 4.23: Table of SNR for images of Figures 4.14 to 4.22. NL-means, TV- ϵ and TV-LSE share the largest SNR values. Now visually NL-means is better for large level of denoising whereas SNR values pretend otherwise; TV-LSE is visually better than TV- ϵ for middle level of denoising. SNR is far from being a perfect indicator of image quality.

Experiments were also carried out on a pure noise image. Pure noise images imitate regions of originally smooth images. A good denoising method is expected to reduce the noise without creating structure. This is a fundamental requirement in applications such as satellite image interpretation, and more generally when denoising is a prior step to feature detection. Figure 4.24 shows the denoised images obtained by each of the previous methods, for a given level of denoising. The results can be classified into 4 main behaviors:

- images with turbulent texture created: information has been created and this leads to highly unnatural images (Mean Curvature Motion and Gaussian convolution);
- images which are not far from being piecewise constant: this is the staircasing effect with spurious contrasted edges (TV-MAP and TV- ϵ);
- images with a cloudy texture, like a colored noise (TV-LSE and Wiener filter);
- images containing a much reduced noise linearly combined with low-frequency texture (TV-barycentre and NL-means).

It is clear that the last two cases are more desirable than the two first because the corresponding images are far less informative and contain no significant features.




Figure 4.24: An i.i.d. white noise uniformly drawn on $[0, 255]$ (top-left image) is denoised by different methods, maintaining a constant level of denoising $\|\hat{u} - v\| = 68.5$. A unique affine contrast change has been processed on all the images of the figure in order to improve the visibility of details. This is why some gray levels are saturated. The TV-LSE denoised image (top-center, with parameters $(\lambda, \sigma) = (150, 20)$) is more regular than the original but few structure is created. TV-MAP denoising (top-right, with $\lambda_{MAP} = 104$) creates big artificial structures due to staircasing effect. In the TV-barycentre denoised image (center-left, with $(\lambda, t) = (150, 0.943)$), the added noise hides the TV-MAP staircasing effect, structure is replaced by some noise. The TV- ϵ denoised image (center, with $(\epsilon, \lambda_\epsilon) = (1, 109.5)$) looks quite the same as TV-MAP. The Wiener filter (center right, with $\lambda_{Wiener} = 17.53$) looks like i.i.d. colored noise. No particular structure is created. The result is similar to TV-LSE. The Mean Curvature Motion (bottom-left, with $T = 2.33$) gives rise to a structured turbulent texture. A Gaussian convolution (bottom center, with $g = 1.38$) creates smooth structures. The NL-means (bottom-right, with $h = 61.75$) is an almost perfect white noise with reduced s.d., although low frequency texture appears.

4.4 Experiments on synthetic signals

In this section experiments are carried out on synthetic signals. The transposition of the algorithm from images to signals is straightforward, with the standard numerical scheme for total variation

$$TV(u) = \sum_{x=1}^n |u(x+1) - u(x)|$$

where n is the size of the signal. The precision parameter ε in the stopping criterion should be taken less than 1 because the error made by $\varepsilon \approx 1$ is still visible. In practice $\varepsilon \approx 0.3$ is enough to have a good visual quality.

Figures 4.25 and 4.26 come back to the synthetic signals of Figure 1.3 in Chapter 1, and compare TV-MAP and TV-LSE behavior on a Sine curve borne by an affine ramp and on a noisy ramp. In the two cases, TV-LSE shrinks the noise amplitude and preserves the global Sine shape. In other words, TV-LSE respects the geometry of the signal while TV-MAP introduces staircasing and breaks the regularity of the original signal.

Figure 4.27 compares a chirp (oscillating signal with increasing frequency) denoised by TV-LSE and TV-MAP. In the two cases the oscillations are shrunk with respect to their frequencies, but the Sine shape is only preserved by TV-LSE while TV-MAP decimates the signal's extrema (again a staircasing effect) and creates singularities.

Figure 4.28 shows the denoising of an isolated peak on a null signal degraded by a Gaussian white noise. For TV-MAP as well as for TV-LSE, the peak's amplitude shrinks when the level of denoising increases, and the peak contaminates its neighborhood with positive values. The TV-LSE signal contains noise even in high levels of denoising (this is compatible with the invertibility of the LSE operator (see Chapter 3 Subsection 3.1.3)), but contrary to TV-MAP no particular feature (such as edges) is created.

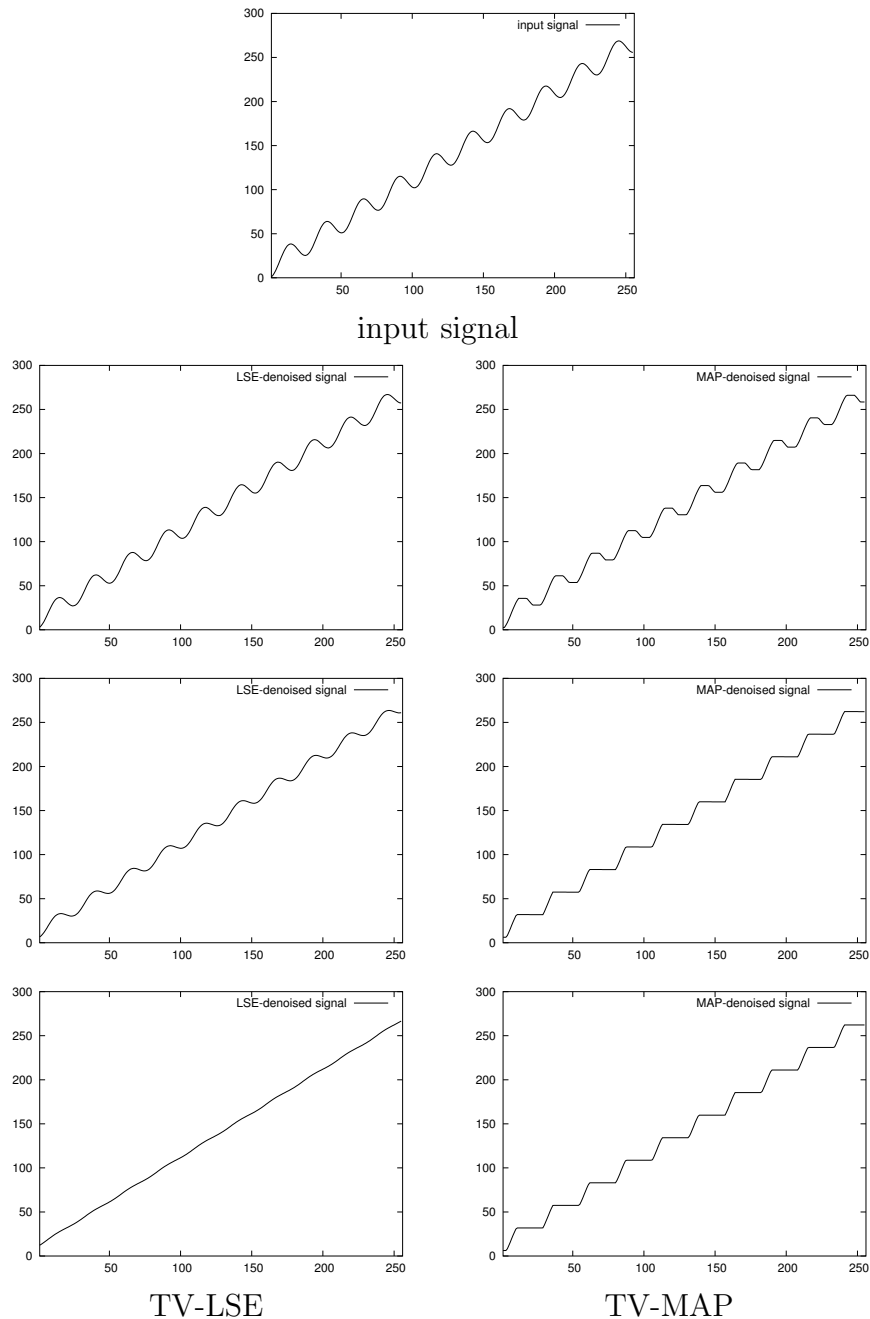


Figure 4.25: Experiments on a synthetic signal – ramp modulated by a Sine function. In the TV-LSE denoised signals (left column, from top to bottom: $(\lambda, \sigma) = (80, 20)$, $(\lambda, \sigma) = (200, 20)$, $(\lambda, \sigma) = (2000, 20)$) the amplitude of the oscillations decreases when the level of denoising increases, but the Sine shape is preserved. On the other hand the TV-MAP denoised signals (right column, from top to bottom: $\lambda_{MAP} = 10.4$, $\lambda_{MAP} = 38.8$, $\lambda_{MAP} = 2000$, tuned to have a level of denoising equal to the one of the TV-LSE denoised signal it faces, except for the bottom image for which the level of denoising cannot be reached because the total variation cannot be decreased anymore) start by having their local maxima and minima saturated; then for high level of denoising the Sine shape of the input totally disappears and staircasing effect is ubiquitous. The algorithms have been run on a much longer signal (three times longer) and the outputs have been cropped, in order to attenuate the boundary conditions effects.

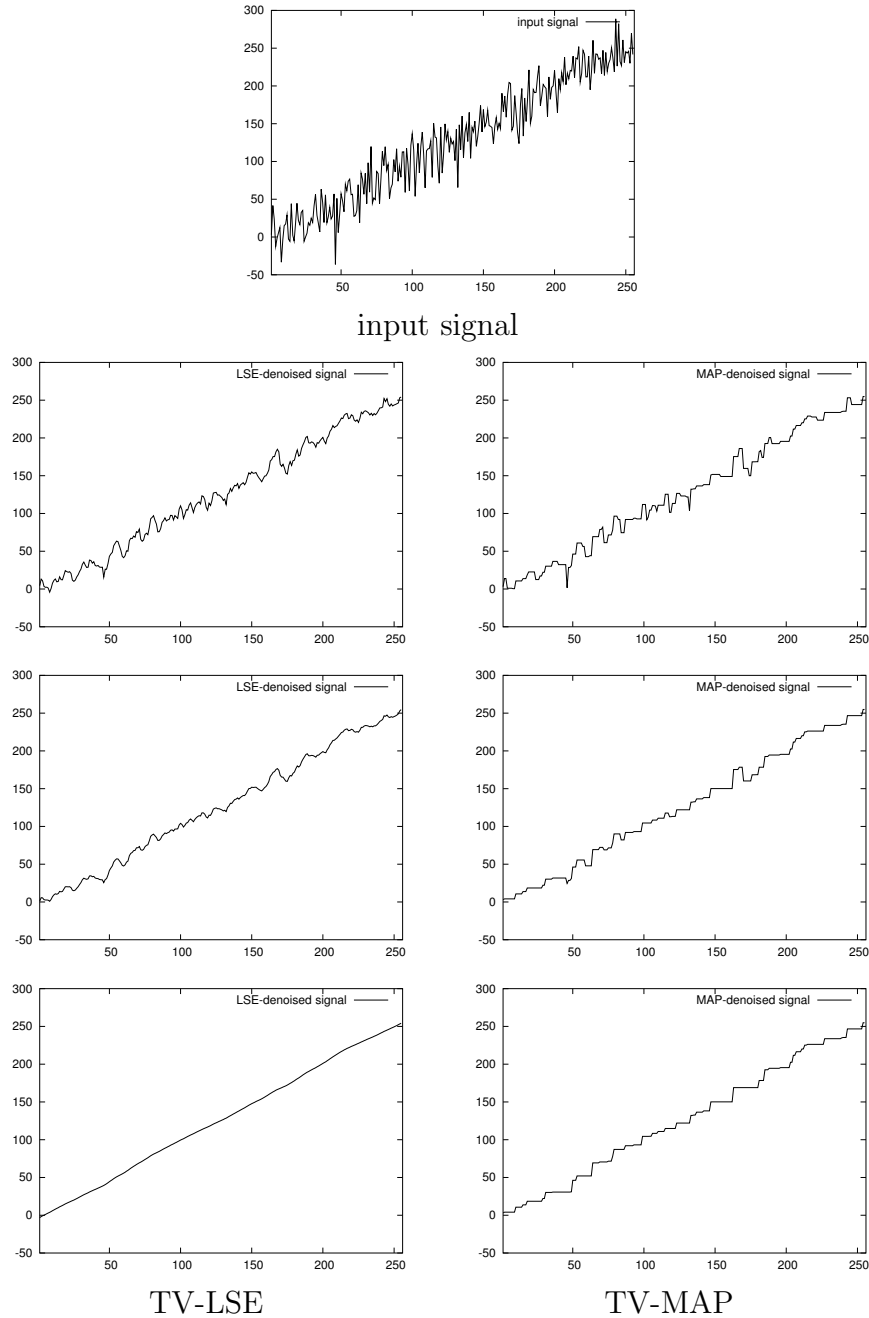


Figure 4.26: Experiments on a synthetic signal – the noisy ramp ($\sigma = 20$). In the TV-LSE denoised signals (left column, from top to bottom: $(\lambda, \sigma) = (80, 20)$, $(\lambda, \sigma) = (150, 20)$, $(\lambda, \sigma) = (2000, 20)$) the signal keeps looking like a noisy ramp, whose noise's variance decreases when the level of denoising increases. Conversely the TV-MAP denoised signals (right column, from top to bottom: $\lambda_{MAP} = 37.9$, $\lambda_{MAP} = 60.8$, $\lambda_{MAP} = 2000$, tuned to have a level of denoising equal to the one of the TV-LSE denoised signal it faces, except for the bottom image for which the level of denoising cannot be reached because the total variation cannot be decreased anymore) present some structural artefacts, with constant intervals separated by sharp discontinuities. This staircasing artefact is emphasized for high levels of denoising: the signal becomes increasing and has longer constant intervals combined with sharper edges. The algorithms have been run on a much longer signal (three times longer) and the outputs have been cropped, in order to attenuate the boundary conditions effects.

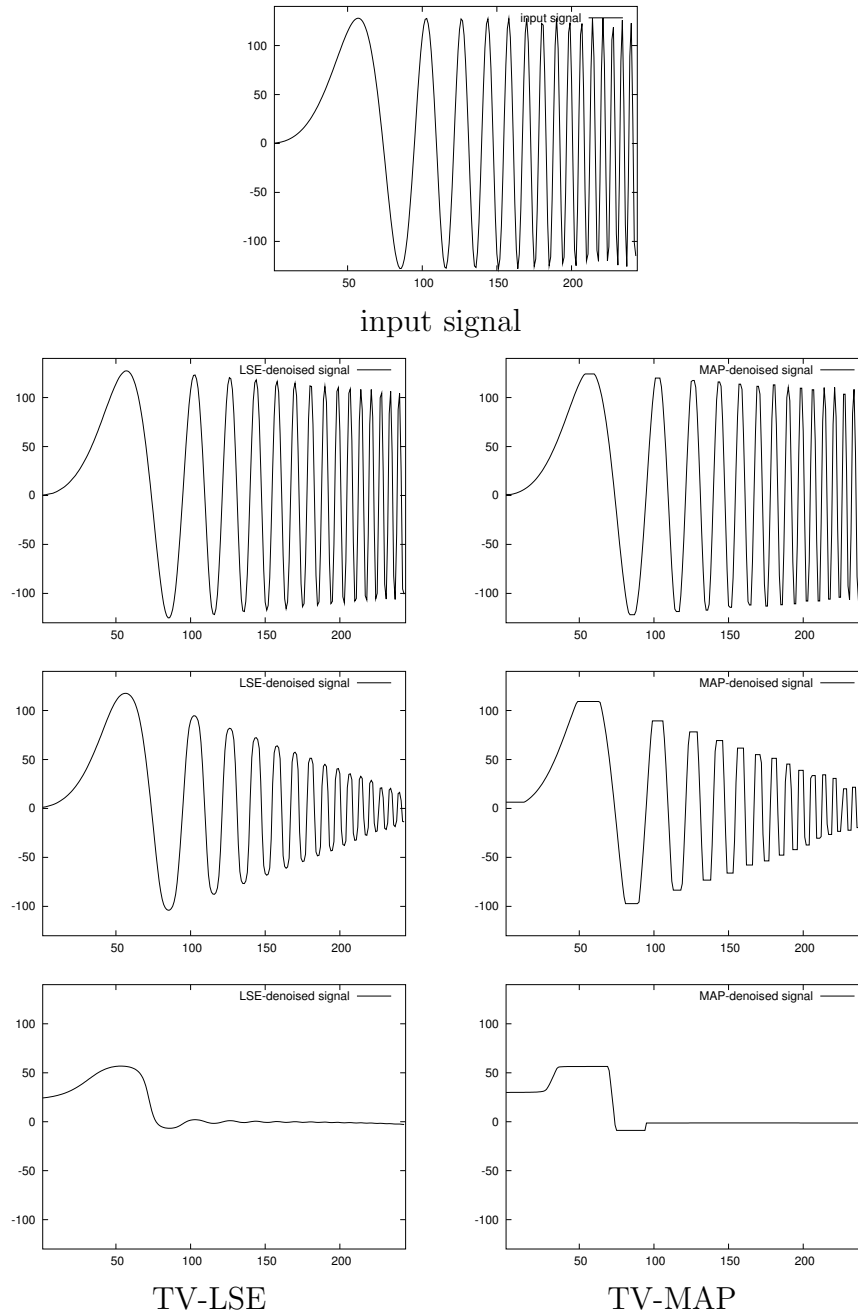


Figure 4.27: Experiments on a synthetic signal – the chirp ($\sin t^3$). In the TV-LSE denoised signals (left column, from top to bottom: $(\lambda, \sigma) = (20, 10)$, $(\lambda, \sigma) = (200, 15)$, $(\lambda, \sigma) = (2000, 20)$) the amplitude of oscillations goes down when the frequency increases and when the level of denoising increases, but the global Sine shape is preserved, until oscillations are totally flattened. In the TV-MAP denoised signals (right column, from top to bottom: $\lambda = 17.49$, $\lambda = 188.8$, $\lambda = 1514$, tuned to have a level of denoising equal to the one of the TV-LSE denoised signal it faces) the amplitude of oscillations also goes down when the frequency increases and when the level of denoising increases, but contrary to TV-LSE denoised signals, staircasing appears again at any frequency, thus breaking the smoothness of the input signal and creating artificial structure.

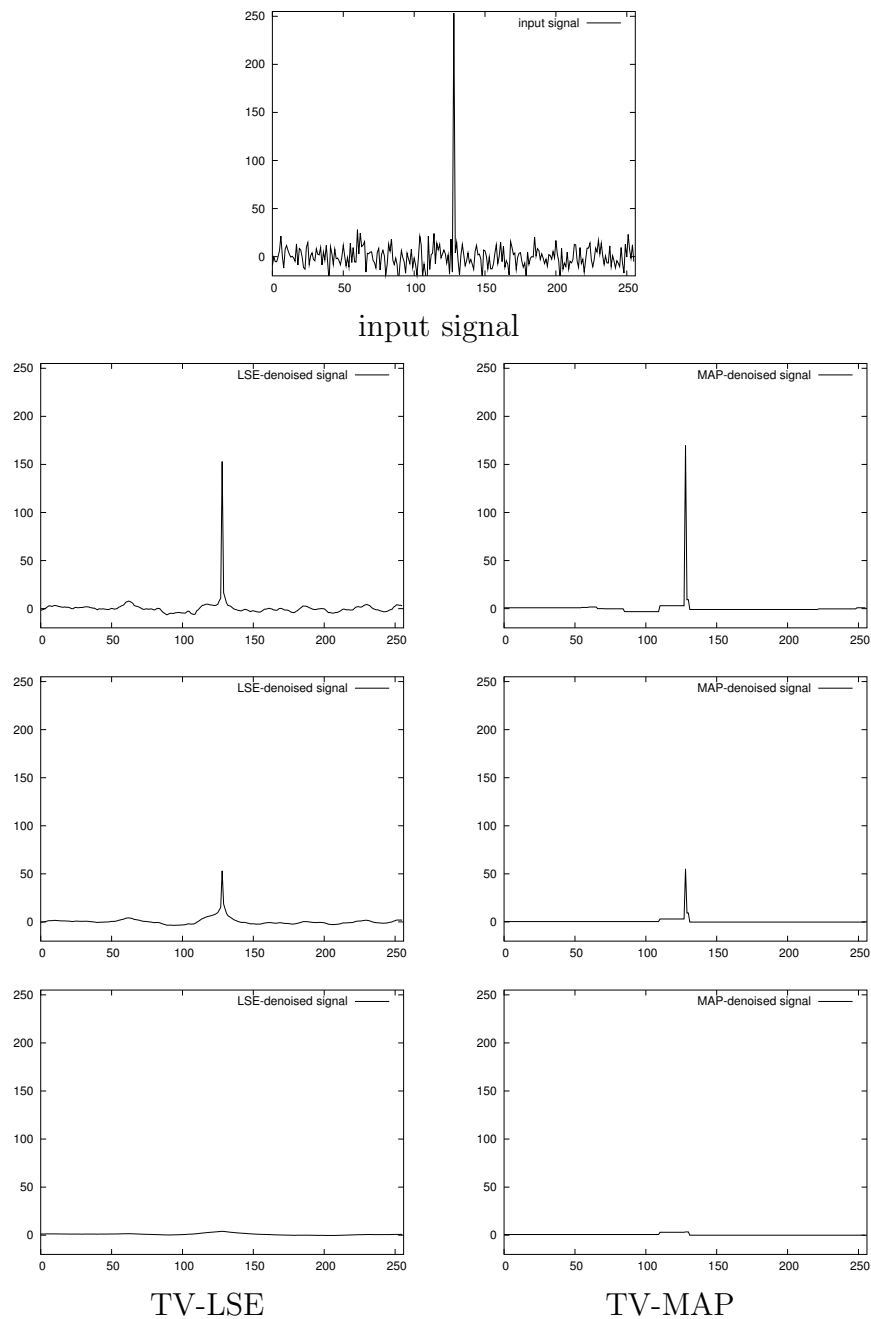


Figure 4.28: Experiments on a synthetic signal – the noisy peak. In the TV-LSE denoised signals (left column, from top to bottom: $(\lambda, \sigma) = (100, 10)$, $(\lambda, \sigma) = (200, 10)$, $(\lambda, \sigma) = (800, 10)$) In the TV-MAP denoised signals (right column, from top to bottom: $\lambda = 83.5$, $\lambda = 198.0$, $\lambda = 261.4$, tuned to have a level of denoising equal to the one of the TV-LSE denoised signal it faces)

4.5 Staircasing and isolated pixels artefacts

In this section we focus on the particular artefacts encountered by TV-MAP denoising, namely the staircasing effect and the emergence of contrasted isolated pixels.

First, concerning the staircasing effect, Figure 4.29 comes back to the images of Figures 1.4 and 1.5 (from Chapter 1) and compares TV-MAP and TV-LSE denoising with equal level of denoising. Experiments agree with the theorem stating the absence of staircasing in TV-LSE denoised images. In TV-MAP denoising, staircasing effect mostly occurs on regions where the noise or small scale texture dominates the geometry of the image, and corrupts smooth zones into piecewise constant zones with spurious contrasted edges. In TV-LSE denoising such an artefact does not appear, sometimes at the expense of a noisier aspect.

As seen in previous experiments, taking a very small σ can lead to visual staircasing effect, for the image is very close to the TV-MAP denoised image. In all the experiments of this section σ was chosen equal to the standard deviation of the noise, and this value appeared to be large enough to avoid a visual staircasing effect.

Figures 4.30 and 4.31 deal with the isolated pixels generated by TV-MAP denoising. This isolated pixel artefact occurs when the noise dominates the geometry of the image, and when the level of denoising is weak. In this case a good denoising operator is expected to shrink amplitude of the noise with no creation of artificial features, which is not the case with TV-MAP. The ideal case is when a constant original image is degraded by noise. Figure 4.30 shows the denoising of such a synthetic Gaussian noise image. We check that the isolated pixels generated by TV-MAP vanish when denoising via TV-LSE. Figure 4.31 shows the same behavior on a natural image (Barbara image). Several sets of parameters (λ, σ) achieving some desired level of denoising make TV-LSE greatly reduce the visual impact of the isolated pixels generated by TV-MAP.



Figure 4.29: Reduction of the staircasing effect. We take the images of Figures 1.4 and 1.5, and experimentally check that TV-LSE does not create staircasing artefacts. The left column shows zooms of Lena and Barbara images with Gaussian white noise added (s.d. 10). The middle column shows the corresponding TV-LSE denoised images $((\lambda, \sigma) = (40, 10))$, while the right column shows the TV-MAP denoised images corresponding to equivalent level of denoising (from top to bottom: $\lambda_{MAP} = 25.6$, $\lambda_{MAP} = 20.3$, $\lambda_{MAP} = 29.0$, $\lambda_{MAP} = 26.9$). Staircasing effect is huge in TV-MAP but does not occur in TV-LSE.




Figure 4.30: Isolated pixels on noise. A Gaussian white noise (with s.d. 40, left image) is denoised using TV-LSE (with $(\lambda, \sigma) = (80, 40)$, center image) and TV-MAP (with $\lambda_{MAP} = 35.9$, right image). The parameters have been chosen to give equal levels of denoising. TV-LSE yields in a correlated noise-like image, with hardly no structure created, while TV-MAP creates highly contrasted isolated pixels and staircasing. Contrast has been enhanced (a unique affine contrast change for the 3 images) in order to improve the visualization of the denoised images.




Figure 4.31: Isolated pixels on a natural image. A detail of Barbara has been noised by a Gaussian white noise with s.d. 10 (top-left image). TV-MAP denoising (top-right image, with $\lambda_{MAP} = 8$) is applied to the noisy image, as well as TV-LSE denoising (bottom-left image with $(\lambda, \sigma) = (13.1, 10)$, bottom-right image with $(\lambda, \sigma) = (25.5, 25)$, attaining the same level of denoising as TV-MAP). Many contrasted isolated pixels corrupt the TV-MAP denoised image, while the remaining noise of the TV-LSE denoised images is much more homogeneous. This isolated pixel artefact occurs as soon as TV-MAP is applied with a small level of denoising on an homogeneous region.

4.6 Other examples, and natural noise

In this section miscellaneous images denoised by TV-LSE are shown, with a comparison to TV-MAP.

In Figures 4.32 and 4.33, details of the classical images Goldhill and Barbara where Gaussian noise has been added are denoised. TV-LSE always gives more natural and pleasant images.

In Figures 4.34 and 4.35, experiments on natural noise are held. We chose images from traditional photography and digitized by a professional (Figure 4.34) or home-made (Figure 4.35) scan. The noise the scan induces is not far from being white and Gaussian. Actually its distribution seems to be more peaked than a Gaussian distribution. Moreover there is no reason for the noise to be independent of the image. Although the noise model does not fit perfectly well to the framework, the TV-LSE denoised images have a correct and natural aspect. In both TV-LSE and TV-MAP, most noise is removed, but in TV-MAP, thin structures as the chain in the top images of Figure 4.34 are erased and staircasing spoils again the image, while in TV-LSE, some noise remains but staircasing is avoided.




Figure 4.32: Examples on Goldhill. On the first column, details of the degraded Goldhill image are shown (Gaussian white noise with s.d. 10). The corresponding TV-LSE denoised images are shown on the second column ($(\lambda, \sigma) = (25, 10)$), and the TV-MAP denoised images on the third column ($\lambda_{MAP} = 14.45, 18.61, 17.55, 16.55$ corresponding to the level of denoising reached by the LSE images). On the first row the isolated pixels of TV-MAP vanish in TV-LSE, and the bushes texture is more natural in TV-LSE. On the second row contrast is a bit more preserved in TV-MAP but the texture of the wall is more credible on TV-LSE. The roofs' stripes of the third row have a better rendering in TV-LSE. The last row is not that clear: a careful observation shows that TV-MAP contains more contrast and more staircasing, while TV-LSE is more blurry.



Figure 4.33: Examples on Barbara. On the first column, details of the degraded Barbara image are shown (Gaussian white noise with s.d. 10). The corresponding TV-LSE denoised images are shown on the second column ($(\lambda, \sigma) = (20, 10)$), and the TV-MAP denoised images on the third column ($\lambda_{MAP} =$ corresponding to the level of denoising of the LSE images).




Figure 4.34: Denoising of naturally noisy images (scanned silver-based photograph). On the first column details of naturally noisy images are shown. On the second column the corresponding TV-LSE denoised images ($(\lambda, \sigma) = (20, 10)$), and on the third column the TV-MAP denoised images with equal level of denoising ($\lambda_{MAP} = 15.8, 13.35, 11.23$). Here again TV-MAP images are nearly piecewise constant, whereas TV-LSE achieves a more natural rendering.




Figure 4.35: Denoising of naturally noisy images (scanned silver-based photograph). On the first column details of naturally noisy images are shown. On the second column the corresponding TV-LSE denoised images $((\lambda, \sigma) = (20, 10))$, and on the third column the TV-MAP denoised images with equal level of denoising $(\lambda_{MAP} = 15.56, 14.10)$.

Part II

Locality and prior

Chapter 5

Locality in Total Variation denoising

Résumé Dans ce chapitre nous proposons des versions localisées du débruitage par variation totale. Nous en analysons quelques aspects: compromis biais-variance, EDP limite, pondération des voisinages, etc.

Abstract In this chapter we discuss localized versions of Total Variation denoising. Several aspects are analyzed: bias-variance trade-off, limiting PDE, neighborhood weighting, etc.

5.1 Introduction

Total variation is a penalizing quantity which has proven to be quite well adapted to variational image denoising [Rudin et al. 1992]. This model brings into play short-range as well as long-range interactions between the pixels. Indeed, the Total Variation is quantity which gives a measurement of global regularity of images, and in the denoising model the denoised gray level assigned to a pixel generally depends on the whole image. In this global framework, denoising cannot commute with image cropping, because when the image is cropped, the information formerly brought by far away pixel is lost and the denoising may lead to a different output. For natural images where two regions of the image may come from very different scene situations, it can sound weird that the denoising of one pixel needs the contribution of remote pixels. One should rather expect that the denoising of a pixel only needs the information brought by its neighbors, provided that the neighborhood is large enough.

This idea of localization of denoising has already been exploited in a completely different framework, that is in the Non-Local means (NL-means) algorithm [Buades et al. 2006a]. The principle of this algorithm is to make all the pixels contribute to the denoising of one of them, by assigning to it a weighted average of all the gray values which occur in the image. The weights must be subtly selected, but under a stationary assumption this method is mathematically proven to converge towards the conditional expectation of the original image, when the image's size goes to infinity. However it has been noticed that in practice, the denoising performances were better when considering smaller images. This amounts to constrain the denoising of a pixel to depend on the pixels lying in a certain neighborhood of it only. This localization procedure has been interpreted as [Kervrann and Boulanger 2006] as a bias-variance trade-off.

Despite huge differences separating the NL-means and Total Variation (TV) denoising, localization can be transposed to TV denoising. This chapter is devoted to an (uncomplete) exploration of localized versions of the TV minimization denoising model. The most straightforward local version is simply defined by the minimization of the energy of [Rudin et al. 1992] where the data is “cropped” along the boundary of a window. If discrete images are used, and if this window is denoted \mathcal{W} , this amounts to minimizing

$$\sum_{x \in \mathcal{W}} (u(x) - v(x))^2 + \lambda \sum_{x \in \mathcal{W}} |\nabla u(x)|,$$

where λ is the regularity parameter of the model, which controls the amount of denoising. The summations are taken over the window \mathcal{W} instead of the whole domain, and the associated denoising becomes an instance of a neighborhood filter.

In this chapter two local models deduced from this local energy minimization are developed. First, Section 5.2 gives precision about the global character of the TV denoising. Then in Section 5.3, properties of the local denoising model where the regularity parameter λ is not window-dependent are derived. A model justification is provided by a Bayesian interpretation. Several regularity results are given; in particular a link to the heat equation is heuristically driven. The window size is discussed in terms of bias-variance trade-off. Then in Section 5.4 this regularity parameter λ is made \mathcal{W} -dependent. It is tuned in such a way that the quantity of noise removed from a window \mathcal{W} is a user-given constant. This approach is linked to the constrained Total Variation model originally proposed by Rudin et al. [1992], formerly to the Lagrangian approach: the noise’s standard deviation (s.d.) σ is known, and the Total Variation is minimized over the ball with radius σ centered on the noisy image. In both sections numerical experiments are carried out. Concluding remarks end the chapter at Section 5.5.

5.2 Total variation minimization is global

The Total Variation minimization proposed by [Rudin et al. 1992], takes a noisy image v and gives back the image which minimizes an energy written

$$\mathcal{E}(u) = \|u - v\|^2 + \lambda TV(u).$$

As seen in Chapter 1, the image which minimizes \mathcal{E} can be interpreted as a Maximum A Posteriori estimate of the original image, assuming a white Gaussian model for the noise

$$p_n(n) = \frac{1}{Z} e^{-\frac{\|n\|^2}{2\sigma^2}}$$

and the Total variation distribution

$$p(u) = \frac{1}{Z} e^{-\beta TV(u)}$$

as a prior on images, provided that $\lambda = 2\sigma^2\beta$ (Z is a “universal” normalizing constant which makes the equations simpler).

This Total Variation prior distribution is a Gibbs model for images. Indeed a scheme for the gradient norm needs the selection of a neighborhood shape \mathcal{N} which will be interpreted as a clique in the Gibbs field

$$\frac{1}{Z} e^{-\beta \sum_{x \in \Omega} f(u(\mathcal{N}_x))}$$

where f is the potential function associated to the clique \mathcal{N} . For example, the gradient scheme defined by

$$|\nabla u(i, j)| = \sqrt{(u(i+1, j) - u(i, j))^2 + (u(i, j+1) - u(i, j))^2}$$

is associated to a neighborhood \mathcal{N}_x containing the pixel x , its top and its right neighbors. As in any Gibbs field (Hammersley-Clifford theorem [Hammersley and Clifford 1971; Besag 1974; Chalmond 2003]), the gray level at a pixel of a random image distributed as the TV prior only depends on the gray levels located in a slightly larger neighborhood, as is shown by

$$\begin{aligned} p(u(x) | (u(y))_{y \neq x}) &= \frac{1}{Z} e^{-\beta \sum_{y \in \Omega} (f(u(\mathcal{N}_y)) - f(u_{/x}(\mathcal{N}_y)))} \\ &= \frac{1}{Z} e^{-\beta \sum_{y \in \mathcal{N}_x} (f(u(\mathcal{N}_y)) - f(u_{/x}(\mathcal{N}_y)))} \\ &= p(u(x) | u(\mathcal{N}_x^2 \setminus \{x\})) \end{aligned}$$

where

$$u_{/x}(y) = \begin{cases} u(y) & \text{if } y \neq x \\ 0 & \text{if } y = x, \end{cases} \quad \text{and} \quad \mathcal{N}_x^2 = \{y \mid \exists x_0 \in \mathcal{N}_x, y \in \mathcal{N}_{x_0}\}.$$

Despite the assumption of a white Gaussian noise model, still, it is well known the local interactions propagate step by step and the posterior distribution of the gray level of one pixel x alone depends on the whole image v [Besag 1986]. Let us give the example of a Dirac image, which shows that the denoising is global indeed. Let v be an image defined on a discrete domain Ω by

$$v(x) = \begin{cases} A & \text{if } x = x_0 \\ 0 & \text{else.} \end{cases}$$

Then there exists $\alpha > 0$ (depending on λ and on the position of x_0 in the domain Ω) such that the globally denoised version of v satisfies

$$Tv(x) = \begin{cases} \begin{cases} A - \alpha & \text{if } x = x_0 \\ \frac{\alpha}{|\Omega| - 1} & \text{else,} \end{cases} & \text{when } A > \alpha, \\ Tv(x) = \frac{A}{|\Omega|} \text{ for all } x & \text{otherwise.} \end{cases}$$

This proves that any pixel in Ω is influenced by the gray level at x_0 , even if x and x_0 are separated by a large distance. Examples of global denoising for images defined on a continuous domain are given in Strong and Chan [1996].

5.3 Lagrangian local TV-MAP

5.3.1 Construction

In this section the total variation denoising is compelled to become local. For a given window shape $\mathcal{W} \subset \mathbb{Z}^2$ containing 0, generally larger than the neighborhood \mathcal{N} of the gradient norm scheme, we define the window \mathcal{W}_x defined by the translation of \mathcal{W} with vector x

$$\mathcal{W}_x = (x + \mathcal{W}) \cap \Omega,$$

where Ω stands for the domain where the images are defined. Let $x \in \Omega$ be a pixel. A local prior on images defined on \mathcal{W}_x is defined by

$$p(u(\mathcal{W}_x)) = \frac{1}{Z} e^{-\beta TV(u(\mathcal{W}_x))},$$

together with a local Gaussian model on the noise

$$p_{n(\mathcal{W}_x)}(n(\mathcal{W}_x)) = \frac{1}{Z} e^{-\frac{\|n(\mathcal{W}_x)\|^2}{2\sigma^2}}.$$

Such distributions lead to the posterior distribution of the window $u(\mathcal{W}_x)$ conditionally to $v(\mathcal{W}_x)$

$$p(u(\mathcal{W}_x)|v(\mathcal{W}_x)) = \frac{1}{Z} e^{-\frac{\|u(\mathcal{W}_x) - v(\mathcal{W}_x)\|^2}{2\sigma^2}} \cdot e^{-\beta TV(u(\mathcal{W}_x))}.$$

Any estimate of $u(\mathcal{W}_x)$ based on this posterior distribution $p(u(\mathcal{W}_x)|v(\mathcal{W}_x))$ is compelled to depend on the noisy window $v(\mathcal{W}_x)$ only. In this section we focus on the Maximum A Posteriori estimate of the window $u(\mathcal{W}_x)$, which is given by

$$\widehat{u(\mathcal{W}_x)} = \arg \min_{w \in \mathbb{R}^{\mathcal{W}}} \|w - v(\mathcal{W}_x)\|^2 + \lambda TV(w).$$

The superimposition of windows implies that each pixel has several estimators, one for each window which it belongs to. Namely, for every $h \in \mathcal{W}$, we can get an estimate $\hat{u}_h(x)$ of the gray level $u(x)$, given by

$$\hat{u}_h(x) = w(h) \quad \text{where} \quad w = \arg \min_{w \in \mathbb{R}^{\mathcal{W}}} \|w - v(\mathcal{W}_{x-h})\|^2 + \lambda TV(w).$$

A smart aggregation of these estimators would probably lead to a more powerful estimator, but we choose here to consider the simplest aggregate, i.e. the estimate corresponding to $h = 0$, writing

$$T_{\mathcal{W}}v(x) = w(0) \quad \text{where} \quad w = \arg \min_{w \in \mathbb{R}^{\mathcal{W}}} \|w - v(\mathcal{W}_x)\|^2 + \lambda TV(w). \quad (5.1)$$

This model where the regularity parameter λ does not depend on the location x is consistent with a prior independent of x too. In the following the shape \mathcal{W} is either a square or a disc, and a pixel x denoised by $T_{\mathcal{W}}$ depends on $v(y)$ with $|y - x|_{\infty}$ or $|y - x|_2$ less than a constant characterizing the size of \mathcal{W} .

5.3.2 Algorithm

A very simple algorithm consists in applying Chambolle’s dual algorithm [Chambolle 2004] on every window $v(\mathcal{W}_x)$ of the noisy image to minimize the corresponding local energy (5.1).

However several tests of this algorithm running on signals and images reveal an annoying artefact: structures with a given frequency can be shrunk to zero while others with higher frequency can come out again. That is why it can be considered as an *aliasing* artefact. It occurs when the signal or image contains structure with very high frequencies and when the window and the parameter λ are very large. It is illustrated on Figure 5.1 with the chirp. It can also be detected on images but it is much harder to see.




Figure 5.1: The effect of local denoising on the chirp. The chirp (left) is denoised using the local TV-MAP denoising with first a 21×21 window (center), and then a 61×61 window (right), with the same $\lambda = 2000$). We would expect to obtain the same shape, with decreasing amplitude in the oscillations (high frequencies would be considered as noise). However severe aliasing artefacts are visible in the two signals. For instance, on the central signal, rebounds of the signal can be seen in the high frequencies, and in the signal on the right even the shape of the first arch is ruined.

This phenomenon is due to the boundary conditions which induce distortions on the denoising, even though we chose free boundary conditions on the windows (as Chambolle’s algorithm tells). Thus we try to “smooth” the window as in the NL-means algorithm [Buades et al. 2005]. Namely, instead of considering a space-homogeneous data-fidelity term, we introduce a weighted norm $\|\cdot\|_{2,\omega}$ defined on $\mathbb{R}^{\mathcal{W}}$ by

$$\|u\|_{2,\omega}^2 = \sum_{(i,j) \in \mathcal{W}} \omega_{ij} |u_{ij}|^2, \quad \text{where } \forall (i,j) \in \mathcal{W}, \omega_{ij} > 0 \text{ and } \sum_{(i,j) \in \mathcal{W}} \omega_{ij} = |\mathcal{W}|.$$

Typically the weights (ω_k) can be taken Gaussian, i.e.

$$\omega_{ij} = \frac{|\mathcal{W}|}{\sum_{(i,j) \in \mathcal{W}} e^{-\frac{i^2+j^2}{2a^2}}} e^{-\frac{i^2+j^2}{2a^2}}. \quad (5.2)$$

Notice that the limiting case where $a \rightarrow +\infty$ reduces to the classical Euclidean norm on the window. The algorithm described hereafter is suitable for any positive weights, but is useful in reducing the artefacts only if the outer pixels have lower weights than inner pixels.

Let $x \in \Omega$. When considering the minimization of the smoothed energy

$$\mathcal{E}_\omega(p) = \|p - v(\mathcal{W}_x)\|_{2,\omega}^2 + \lambda TV(p) \quad (5.3)$$

on a window \mathcal{W}_x , the norm $\|\cdot\|_{2,\omega}$ reduces the impact of the border pixels, and hence makes the boundary effects less effective.

Let us derive the ideas of Chambolle's algorithm to this energy. For that, let us introduce the invertible diagonal operator D such that

$$\forall u \in \mathbb{R}^{\mathcal{W}}, \forall k \in \mathcal{W}, (Du)_k = \omega_k u_k$$

(D is always invertible since $\omega_k > 0$, and converges to the identity when $a \rightarrow \infty$ for Gaussian weights (5.2)). Notice that for any image u , we have $\|u\|_{2,\omega} = \|D^{1/2}u\|$.

We also assume as in [Chambolle 2004] that the domain of the window \mathcal{W} is a square $N \times N$, and that the total variation is computed from the gradient discretized as follows

$$(\nabla u)_{i,j}^1 = \begin{cases} u_{i+1,j} - u_{i,j} & \text{if } i < N \\ 0 & \text{if } i = N \end{cases} \quad \text{and} \quad (\nabla u)_{i,j}^2 = \begin{cases} u_{i,j+1} - u_{i,j} & \text{if } j < N \\ 0 & \text{if } j = N, \end{cases}$$

and that the norm of the gradient is L^2 on its two components. The divergence operator dual to the gradient is then defined on $(\mathbb{R}^{\mathcal{W}})^2$ by

$$(\operatorname{div} p)_{i,j} = \begin{cases} p_{i,j}^1 - p_{i-1,j}^1 & \text{if } 1 < i < N \\ p_{i,j}^1 & \text{if } i = 1 \\ -p_{i-1,j}^1 & \text{if } i = N \end{cases} + \begin{cases} p_{i,j}^2 - p_{i,j-1}^2 & \text{if } 1 < j < N \\ p_{i,j}^2 & \text{if } j = 1 \\ -p_{i,j-1}^2 & \text{if } j = N. \end{cases}$$

Then the next proposition holds.

Proposition 5.1 *Under the above conditions, the minimizer of the smoothed energy \mathcal{E}_ω (5.3) writes*

$$T_\omega v(\mathcal{W}_x) = v(\mathcal{W}_x) - \pi_{\frac{\lambda}{2}K}(v(\mathcal{W}_x))$$

where

$$K = \{D^{-1}(\operatorname{div} p) \mid p \in \mathbb{R}^{\mathcal{W}}, \forall (i,j) \in \mathcal{W}, |p_{i,j}| \leq 1\}$$

is closed and convex, and $\pi_{\frac{\lambda}{2}K}$ denotes the projection operator on $\frac{\lambda}{2}K$. Furthermore, letting $\pi_{\frac{\lambda}{2}K}(v(\mathcal{W}_x)) = \frac{\lambda}{2}D^{-1}(\operatorname{div} p)$, p satisfies

$$\forall (i,j) \in \mathcal{W}, |(\nabla(\frac{\lambda}{2}D^{-1}(\operatorname{div} p) - v(\mathcal{W}_x)))_{i,j}| p_{i,j} = (\nabla(\frac{\lambda}{2}D^{-1}(\operatorname{div} p) - v(\mathcal{W}_x)))_{i,j}.$$

Proof : We follow the proof of Chambolle and generalize it to the weighted case. Considering the subgradient of \mathcal{E}_ω (in (5.3)), we obtain that the image $u \in \mathbb{R}^{\mathcal{W}}$ minimizing \mathcal{E}_ω is such that

$$2D(u - v(\mathcal{W}_x)) + \lambda \partial TV(u) \ni 0,$$

i.e.

$$\frac{Dv(\mathcal{W}_x) - Du}{\lambda/2} \in \partial TV(u).$$

Now denoting TV^* the Legendre-Fenchel transform of TV , we get thanks to the Moreau decomposition (see [Moreau 1962] or [Hiriart-Urruty and Lemaréchal 1993])

$$u \in \partial TV^*\left(\frac{Dv(\mathcal{W}_x) - Du}{\lambda/2}\right).$$

Now denoting $w = \frac{v(\mathcal{W}_x) - u}{\lambda/2}$, we get

$$w - \frac{v(\mathcal{W}_x)}{\lambda/2} + \frac{2}{\lambda} \partial TV^*(Dw) \ni 0,$$

whose left-hand term is the subgradient of the energy $\frac{1}{2}\|w - \frac{v(\mathcal{W}_x)}{\lambda/2}\|^2 + \frac{2}{\lambda} TV^*(Dw)$. It means that w minimizes the latter energy. But TV^* is the characteristic function of a convex set [Chambolle 2004], and in particular

$$TV^*(Dw) = \begin{cases} 0 & \text{if } w \in K \\ +\infty & \text{else,} \end{cases} \quad \text{where } K = \{w = D^{-1} \operatorname{div} p \mid \forall (i, j) \in \mathcal{W}, |p_{i,j}| \leq 1\},$$

since D is invertible. K is closed and convex because $p \mapsto D^{-1}p$ is linear. Hence the minimization of the energy $\frac{1}{2}\|w - \frac{v(\mathcal{W}_x)}{\lambda/2}\|^2 + TV^*(Dw)$ amounts to project $v(\mathcal{W}_x)$ onto $\frac{\lambda}{2}K$. In other words,

$$w = \arg \min \left\{ \|v(\mathcal{W}_x) - \frac{\lambda}{2} D^{-1} \operatorname{div} p\|^2 \mid w = D^{-1} \operatorname{div} p, \forall (i, j) \in \mathcal{W}, |p_{i,j}| \leq 1 \right\}. \quad (5.4)$$

The necessary and sufficient Karush-Kuhn-Tucker conditions hold and there exists $\alpha \in \mathbb{R}^{\mathcal{W}}$ such that

$$\forall (i, j) \in \mathcal{W}, -\nabla \left(\frac{\lambda}{2} D^{-1} \operatorname{div} p - v(\mathcal{W}_x) \right)_{i,j} = \alpha_{i,j} p_{i,j} \quad (5.5)$$

with $\alpha_{i,j} \geq 0$ and $\alpha_{i,j}(p_{i,j}^2 - 1) = 0$ for all $(i, j) \in \mathcal{W}$. Then either $\alpha_{i,j} > 0$ and $|p_{i,j}| = 1$, or $|p_{i,j}| < 1$ and $\alpha_{i,j} = 0$. In both cases $\alpha_{i,j} = |\nabla \left(\frac{\lambda}{2} D^{-1} \operatorname{div} p - v(\mathcal{W}_x) \right)_{i,j}|$, and replacing this value of $\alpha_{i,j}$ in (5.5) concludes the proof. \square

Now following Chambolle [2004], a semi-implicit gradient descent with step $\tau > 0$ given by

$$p_{i,j}^{n+1} = p_{i,j}^n + \tau \left(\nabla(D^{-1} \operatorname{div} p^n - \frac{v(\mathcal{W}_x)}{\lambda/2}) \right)_{i,j} - \left| \nabla(D^{-1} \operatorname{div} p^n - \frac{v(\mathcal{W}_x)}{\lambda/2}) \right|_{i,j} p_{i,j}^{n+1} \quad (5.6)$$

leads to the iterative scheme

$$p_{i,j}^{n+1} = \frac{p_{i,j}^n + \tau \left(\nabla(D^{-1} \operatorname{div} p^n - \frac{v(\mathcal{W}_x)}{\lambda/2}) \right)_{i,j}}{1 + \tau \left| \nabla(D^{-1} \operatorname{div} p^n - \frac{v(\mathcal{W}_x)}{\lambda/2}) \right|_{i,j}}. \quad (5.7)$$

In the unweighted case where $D = Id$, the iterative scheme (5.7) was proven to converge to $\pi_{\frac{\lambda}{2}K}(v(\mathcal{W}_x))$ as soon as $\tau \leq 1/8$ [Chambolle 2004]. Now in the weighted case, the theoretical maximum step size is smaller than $1/8$ and depends on the weights. The following proposition gives an explicit maximum step size for which the iterative scheme converges, thus generalizing Chambolle's result.

Proposition 5.2 *Let $\tau \leq \tau_{max}$, with τ_{max} defined by*

$$\tau_{max} = \frac{1}{4 \max \left\{ \max_{i,j} \left(\frac{1}{\omega_{i,j}} + \frac{1}{\omega_{i+1,j}} \right), \max_{i,j} \left(\frac{1}{\omega_{i,j}} + \frac{1}{\omega_{i,j+1}} \right) \right\}}. \quad (5.8)$$

Let (p^n) defined recursively by (5.7). Then $\frac{\lambda}{2}D^{-1}(\operatorname{div} p^n)$ converges to $\pi_{\frac{\lambda}{2}K}(v(\mathcal{W}_x))$.

Proof : We follow the proof of [Chambolle 2004, Theorem 3.1]. For convenience, let L denote the linear operator $L = D^{-1/2}\operatorname{div}$, and $L^* = -\nabla D^{-1/2}$ the dual operator. First, notice that by induction on n in (5.7), it is straightforward that $|p_{i,j}^n| \leq 1$ for all n . Now let us prove that the sequence $\|Lp^n - \frac{v(\mathcal{W}_x)}{\lambda/2}\|^2$ is decreasing for a certain range of values for τ . Let us set $\eta = (p^{n+1} - p^n)/\tau$. Then the difference of norms

$$\|Lp^{n+1} - \frac{v(\mathcal{W}_x)}{\lambda/2}\|^2 - \|Lp^n - \frac{v(\mathcal{W}_x)}{\lambda/2}\|^2 \quad (5.9)$$

$$\begin{aligned} &= 2\tau \langle L\eta, Lp^n - \frac{v(\mathcal{W}_x)}{\lambda/2} \rangle + \tau^2 \|L\eta\|^2 \\ &= -\tau \left[2\langle \eta, -L^*(Lp^n - \frac{v(\mathcal{W}_x)}{\lambda/2}) \rangle - \tau \|L\eta\|^2 \right] \end{aligned} \quad (5.10)$$

should be negative. Denoting κ the operator norm of L , the last term $\|L\eta\|^2$ can be bounded from above by $\kappa^2\|\eta\|^2$. An overestimation of κ will be given at the end of the proof.

In order to process the first term $2\langle \eta, -L^*(Lp^n - \frac{v(\mathcal{W}_x)}{\lambda/2}) \rangle$ in (5.10), first notice that $\eta = (p^{n+1} - p^n)/\tau$ satisfies, thanks to (5.6),

$$\eta = -L^*(Lp^n - \frac{v(\mathcal{W}_x)}{\lambda/2}) - |L^*(Lp^n - \frac{v(\mathcal{W}_x)}{\lambda/2})|p^{n+1},$$

and consequently, for any pixel (i, j) , the splitting

$$\begin{aligned} 2\eta_{i,j} \cdot (-L^*(Lp^n - \frac{v(\mathcal{W}_x)}{\lambda/2}))_{i,j} &= |\eta_{i,j}|^2 + |(L^*(Lp^n - \frac{v(\mathcal{W}_x)}{\lambda/2}))_{i,j}|^2 - |\eta_{i,j} + (L^*(Lp^n - \frac{v(\mathcal{W}_x)}{\lambda/2}))_{i,j}|^2 \\ &= |\eta_{i,j}|^2 + |(L^*(Lp^n - \frac{v(\mathcal{W}_x)}{\lambda/2}))_{i,j}|^2 - |(L^*(Lp^n - \frac{v(\mathcal{W}_x)}{\lambda/2}))_{i,j}|^2 |p_{i,j}^{n+1}|^2, \end{aligned}$$

allows to state that

$$2\langle \eta, -L^*(Lp^n - \frac{v(\mathcal{W}_x)}{\lambda/2}) \rangle \geq \|\eta\|^2, \quad (5.11)$$

because $|p_{i,j}^{n+1}| \leq 1$ for any $(i, j) \in \mathcal{W}_x$.

Now gathering (5.10) and (5.11) yields

$$\|Lp^{n+1} - \frac{v(\mathcal{W}_x)}{\lambda/2}\|^2 - \|Lp^n - \frac{v(\mathcal{W}_x)}{\lambda/2}\|^2 \leq -\tau [(1 - \tau\kappa^2)\|\eta\|^2]$$

which is negative as soon as $\tau < 1/\kappa^2$. This proves that the sequence $\|Lp^n - \frac{v(\mathcal{W}_x)}{\lambda/2}\|^2$ is decreasing, unless $\eta = 0$, which anyway ensures that $p^{n+1} = p^n$. When $\tau = 1/\kappa^2$, the result remains true, because assuming $\|Lp^{n+1} - \frac{v(\mathcal{W}_x)}{\lambda/2}\| = \|Lp^n - \frac{v(\mathcal{W}_x)}{\lambda/2}\|$, we deduce that

$$|(L^*(Lp^n - \frac{v(\mathcal{W}_x)}{\lambda/2}))_{i,j}| |p_{i,j}^{n+1}| = |(L^*(Lp^n - \frac{v(\mathcal{W}_x)}{\lambda/2}))_{i,j}|,$$

so that either $|p_{i,j}^{n+1}| = 1$ or $L^*(Lp^n - \frac{v(\mathcal{W}_x)}{\lambda/2}) = 0$. In any case, it means that $p^{n+1} = p^n$, thanks to (5.7).

From the decrease of $\|Lp^n - \frac{v(\mathcal{W}_x)}{\lambda/2}\|^2$, the convergence of p^n towards the minimizer of (5.4) can directly proven by adapting Chambolle's proof. Indeed, let m be the limit of $\|Lp^n - \frac{v(\mathcal{W}_x)}{\lambda/2}\|$, and \bar{p} be the limit of a converging subsequence (p^{n_k}) of (p^n) . Letting \bar{p}' denote the limit of (p^{n_k+1}) , we have

$$\bar{p}'_{i,j} = \frac{\bar{p}_{i,j} + \tau(L^*(L\bar{p} - \frac{v(\mathcal{W}_x)}{\lambda/2}))_{i,j}}{1 + \tau|L^*(L\bar{p} - \frac{v(\mathcal{W}_x)}{\lambda/2})_{i,j}|},$$

and repeating the former computations, as $m = \|L\bar{p} - \frac{v(\mathcal{W}_x)}{\lambda/2}\| = \|L\bar{p}' - \frac{v(\mathcal{W}_x)}{\lambda/2}\|$, we can see that $\bar{\eta}_{i,j} = (\bar{p}' - \bar{p})/\tau = 0$ for every (i, j) , i.e. $\bar{p} = \bar{p}'$. Thus

$$-L^*(L\bar{p} - \frac{v(\mathcal{W}_x)}{\lambda/2})_{i,j} = |L^*(L\bar{p} - \frac{v(\mathcal{W}_x)}{\lambda/2})_{i,j}| \bar{p}_{i,j}$$

which is the Euler formulation of a solution of the minimization problem (5.4). Hence [Chambolle 2004], $\frac{\lambda}{2}D^{-1}\operatorname{div}\bar{p}$ is the projection $\pi_{\frac{\lambda}{2}K}(v(\mathcal{W}_x))$ required in Proposition 5.1. Since this projection is unique, we deduce that the whole sequence $\frac{\lambda}{2}D^{-1}\operatorname{div}p^n$ goes to the desired projection.

Now let us give an overestimate (as tight as possible) of the norm κ of the operator L . For any η , we have

$$\begin{aligned} \|L\eta\|^2 &= \|D^{-1/2}\operatorname{div}\eta\|^2 = \sum_{i,j} \frac{1}{\omega_{i,j}} (\eta_{i,j}^1 - \eta_{i-1,j}^1 + \eta_{i,j}^2 - \eta_{i,j-1}^2)^2 \\ &\leq \sum_{i,j} \frac{4}{\omega_{i,j}} ((\eta_{i,j}^1)^2 + (\eta_{i-1,j}^1)^2 + (\eta_{i,j}^2)^2 + (\eta_{i,j-1}^2)^2) \\ &\leq 4 \sum_{i,j} \left(\frac{1}{\omega_{i,j}} + \frac{1}{\omega_{i+1,j}} \right) (\eta_{i,j}^1)^2 + 4 \sum_{i,j} \left(\frac{1}{\omega_{i,j}} + \frac{1}{\omega_{i,j+1}} \right) (\eta_{i,j}^2)^2 \\ &\leq 4 \max \left\{ \max_{i,j} \left(\frac{1}{\omega_{i,j}} + \frac{1}{\omega_{i+1,j}} \right), \max_{i,j} \left(\frac{1}{\omega_{i,j}} + \frac{1}{\omega_{i,j+1}} \right) \right\} \|\eta\|^2. \end{aligned} \quad (5.12)$$

This provides the upper bound on κ^2

$$\kappa^2 \leq 4 \max \left\{ \max_{i,j} \left(\frac{1}{\omega_{i,j}} + \frac{1}{\omega_{i+1,j}} \right), \max_{i,j} \left(\frac{1}{\omega_{i,j}} + \frac{1}{\omega_{i,j+1}} \right) \right\}.$$

This bound is probably not optimal, but the example of

$$\eta_{i,j}^1 = \eta_{i,j}^2 = \mathbf{1}_{(i,j)/\omega_{i,j} \text{ minimum and } (i+1,j+1) \in \Omega}$$

shows that the order of magnitude is correct, since for this η

$$\|L\eta\|^2 \approx \max \left\{ \max_{i,j} \left(\frac{1}{\omega_{i,j}} + \frac{1}{\omega_{i+1,j}} \right), \max_{i,j} \left(\frac{1}{\omega_{i,j}} + \frac{1}{\omega_{i,j+1}} \right) \right\} \|\eta\|^2.$$

Finally, taking $\tau \leq \tau_{max}$ with τ_{max} as in (5.8) yields the desired result. \square

Figure 5.2 shows the chirp denoised with smooth versions of TV-MAP local using the algorithm (the weights are chosen Gaussian as in (5.2)). Attention should be paid to the fact that the weighted version achieves lower levels of denoising than the unweighted one when based on the same window size. Indeed Proposition 5.5 in Subsection 5.3.3 will prove that the weighted denoiser with λ_ω is asymptotically equivalent to the unweighted one with λ (when $\lambda \rightarrow 0$) provided that

$$\lambda_\omega = \frac{|\Omega|}{|\mathcal{W}|} \max \omega_{i,j} \lambda,$$

which tells that we need to take $\lambda_\omega > \lambda$ if we desire an equivalent level of denoising. This is why Figure 5.2 shows two solutions: one with the same window size than for the non-smooth local TV-MAP denoising (with window radius $\approx 3a$, the Gaussian bandwidth, for which the Gaussian becomes negligible), and the other with a larger window able to achieve equal denoising level. In both cases the ‘‘aliasing’’ artefact is perfectly removed.

Let us say a word about the suboptimality of τ_{max} (5.8). First note that

$$\frac{1}{8} \min \omega_{i,j} \leq \tau_{max} \leq \frac{1}{4} \min \omega_{i,j}$$

and that $\tau_{max} \leq 1/8$, which tells that admissible step sizes are smaller than in the unweighted case. In the unweighted case, Chambolle [2004] noticed that the practical maximum step size τ_{max}^{pr} was twice the τ_{max} found in the proof. In the case of Gaussian weights, it seems that

$$\tau_{max}^{pr} \in [2\tau_{max}, 4\tau_{max}]$$

The left bound $2\tau_{max}$ corresponds to the case where the Gaussian bandwidth a is much larger than the window size (then the weights are virtually constant), and the right bound $4\tau_{max}$ corresponds to the case where a is much smaller than the window size (the Gaussian is virtually not truncated). An experiment concerning this gap between τ_{max} and τ_{max}^{pr} is illustrated on Figure 5.3.




Figure 5.2: On the left column a “hard” window is used in the local TV-MAP denoising of the chirp, with two different window sizes (21×21 on the top, 61×61 below, with the same $\lambda = 2000$). Severe aliasing artefacts are visible in the two signals. For instance, on the top signal, rebounds of the signal can be seen in the high frequencies, and in the bottom signal even the shape of the first arch is ruined. The second and third columns show the chirp denoised by a local TV-MAP using a smooth window with still $\lambda = 2000$. In the second column the coefficient a of the smooth window (see Equation (5.2)) was tuned such that the weights become negligible on the borders of the hard window of the first column. The level of denoising (i.e. the L^2 -distance between the input and the denoised signals) obtained by this denoising is much smaller than for the first column, which makes it difficult to compare to the “hard-window” denoising. This is why a third column is shown, where this coefficient a is taken such that the denoised signal achieves the same level of denoising than in the first column. The smooth-window algorithm here allows complete removal of the aliasing artefacts.

Remarks Even if the ratio $\tau_{max}^{pr}/\tau_{max}$ is bounded, both of them go to 0 when $\min \omega_{i,j}$ goes to 0. Indeed, when the weights are far from being uniform (for instance in the case of Gaussian weights with a parameter a much smaller than the window size), the maximal step size τ_{max} given in (5.8) becomes extremely small. In practice in this case, the convergence rate of (p^n) is very slow, and the algorithm becomes untractable.

Another drawback of these smooth windows comes from the model. A Bayesian interpretation of the smoothed energy (5.3) would be that the outer pixels of the window contain more noise than the inner pixels. The distribution of the noise does certainly not reflect the physical reality of the noise.

This is why in our experiments on images the “non-smooth” version of local TV-MAP denoising is used, unless the opposite is explicitly specified. The artefacts which may occur are barely visible in the images, while the smooth version is crucial when dealing with signals, for which the artefact is really shocking.




Figure 5.3: Determination of the practical step τ_{max}^{pr} in the smooth-window local TV-MAP scheme. The theoretical bound τ_{max} (5.8) is actually not optimal. In practice much larger step sizes are able to guarantee convergence of the algorithm (5.7). Chambolle [2004] already noticed that in the non-smooth case, the practical maximal step $\tau_{max}^{pr} = 1/4$ was twice the theoretical $\tau_{max} = 1/8$ given in the convergence proof. Here for truncated Gaussian weights, the relative gap between τ_{max}^{pr} and τ_{max} seems to lie in the interval $[2, 4]$. Each graph corresponds to a certain level of window smoothing, and shows the practical τ_{max}^{pr} (bold line) and the theoretical τ_{max} (thin line) with respect to the regularity parameter λ for a given image. The first graph (a) is relative to a case where the weights $(\omega_{i,j})$ are practically constant on the window ($\min \omega_{i,j} \approx 0.998$, $\max \omega_{i,j} \approx 1.002$, which means that the window is virtually unweighted); it comes that a λ -independent τ_{max}^{pr} is approximately twice the theoretical bound (Chambolle’s case). The graph (b) shows a case where the window is a bit smoother ($\min \omega_{i,j} \approx 0.98$, $\max \omega_{i,j} \approx 1.02$); the relative gap is larger than 2. The graph (c) is relative to a smoother window ($\min \omega_{i,j} \approx 0.794$, $\max \omega_{i,j} \approx 1.24$), the graph (d) to a smoother window ($\min \omega_{i,j} \approx 0.369$, $\max \omega_{i,j} \approx 2.13$), and the last graph (e) to a practically Gaussian window ($\min \omega_{i,j} \approx 0.00231$, $\max \omega_{i,j} \approx 8.12$). The relative gap between τ_{max} and τ_{max}^{pr} increases when the “smoothness” (say $\max \omega_{i,j} / \min \omega_{i,j}$) increases, but never exceeds 4. (These graphs were obtained from the denoising of the same noisy image, by observing the convergence of the ROF energy).

5.3.3 Stability properties

In this subsection we gather several results about local TV-MAP denoising. Unfortunately the assumptions about the discretization of total variation are not the same in the different propositions. Each proposition then comes with its assumptions.

First in Proposition 5.3, which is a typical property from neighborhood filters, the total variation is assumed to be writable as in Proposition 3.5 in Chapter 3

$$TV(u) = \sum_{k \in \Omega} \left(\varepsilon_k^p + \|u(\mathcal{N}_k \cap \Omega) - u_k\|_k^p \right)^{1/p}, \quad (5.13)$$

with $\varepsilon \geq 0$ and $p \in [0, \infty]$, but the proposition also holds in the continuous framework.

Proposition 5.3 (Local comparison principle) *For any pixel x , the denoising operator denoted T applied on the image v satisfies*

$$\min_{y \in \mathcal{W}_x} v(y) \leq Tv(x) \leq \max_{y \in \mathcal{W}_x} v(y).$$

Proof : Let $x \in \Omega$. Let $m = \min_{y \in \mathcal{W}_x} v(y)$ and $M = \max_{y \in \mathcal{W}_x} v(y)$. Let q be the patch associated to the denoising of $v(\mathcal{W}_x)$ i.e. such that

$$q = \arg \min_q \|p - v(\mathcal{W}_x)\|^2 + \lambda TV(q)$$

Assume that $\hat{u}(x) = q(0)$ is not inside $[m, M]$. Let then q' be the thresholded patch associated to q defined by

$$q'(k) = \begin{cases} q(k) & \text{if } q(k) \in [m, M], \\ M & \text{if } q(k) > M, \\ m & \text{if } q(k) < m. \end{cases}$$

Then $\|q' - v(\mathcal{W}_x)\|^2 < \|q - v(\mathcal{W}_x)\|^2$ since the gray levels of q outside $[m, M]$ are set nearer to the noised value $v(y)$. Besides $TV(q') \leq TV(q)$ holds under the condition (5.13) because

$$\forall k \in \Omega, \forall l \in \mathcal{N}_k, |q'_l - q'_k| \leq |q_l - q_k|$$

which implies

$$\forall k \in \Omega, \|q'(\mathcal{N}_k \cap \Omega) - q'_k\|_k^p \leq \|q(\mathcal{N}_k \cap \Omega) - q_k\|_k^p$$

and $TV(q') \leq TV(q)$. The inequality also holds when the images are defined on a continuous domain because the coarea formula [Evans and Gariepy 1992] implies

$$TV(q') = \int_m^M \mathcal{H}^1(\partial\{q' \leq \lambda\}) d\lambda = \int_m^M \mathcal{H}^1(\partial\{q \leq \lambda\}) d\lambda \leq \int_{\mathbb{R}} \mathcal{H}^1(\partial\{q \leq \lambda\}) d\lambda = TV(q),$$

where $\mathcal{H}^1(\partial A)$ denotes the Hausdorff measure of the boundary of A . Hence in any of the two cases, $E(q') < E(q)$, which contradicts the minimality of $E(q)$. \square

The next proposition is a monotony property, and needs a very restrictive condition about TV's discretization. We assume for Proposition 5.4 that it writes

$$TV(u) = \sum_{i,j} |u_{i+1,j} - u_{i,j}| + |u_{i,j+1} - u_{i,j}| \quad (5.14)$$

or that it is defined in the continuous domain. It amounts to assume that TV satisfies the coarea formula [Evans and Gariépy 1992; Darbon and Sigelle 2006a]

$$TV(u) = \int_m^M \mathcal{H}^1(\partial\{u \leq \lambda\}) d\lambda$$

where $\mathcal{H}^1(\partial\{u \leq \lambda\})$ denotes the Hausdorff measure of the boundary of the λ -level set of u . In the discrete case, $\mathcal{H}^1(\partial\{u \leq \lambda\})$ is the perimeter of $\{u \leq \lambda\}$, and is measured thanks to cliques of order 2 [Darbon and Sigelle 2006a].

Proposition 5.4 (Global comparison principle) *Assume that the total variation writes as in (5.14) or that the image domain is continuous. Let $v_1 < v_2$ (i.e. for any pixel x , $v_1(x) < v_2(x)$). Then the local TV-MAP denoised images Tv_1 and Tv_2 satisfy $Tv_1 \leq Tv_2$.*

Proof : Let $v_1 < v_2$ and let $x \in \Omega$. Let p_1 and p_2 the patches which minimize the local energies $E_{v_1(\mathcal{W}_x)}$ and $E_{v_2(\mathcal{W}_x)}$. We can apply the monotony property of the global TV-MAP denoiser ([Chambolle 2005] or [Darbon and Sigelle 2005] for the discrete case, and [Alter et al. 2005a] for the continuous case) on $v_1(\mathcal{W}_x) < v_2(\mathcal{W}_x)$ to prove that $p_1 \leq p_2$. Hence $Tv_1(y) = p_1(0) \leq p_2(0) = Tv_2(y)$. \square

This result is a stability property. Indeed, if a noisy image v_2 is squeezed between two regular images v_1 and v_3 , with $v_1 < v_2 < v_3$ for instance, then the denoising of the the 3 images will preserve the order $Tv_1 \leq Tv_2 \leq Tv_3$. If v_1 and v_3 are regular, we can expect to have $Tv_1 \approx v_1$ and $Tv_3 \approx v_3$. Then Tv_2 is expected to be squeezed between v_1 and v_3 .

As noticed in Subsection 5.3.2, the level of denoising of a patch decreases when considering weighted windows instead of unweighted windows. In order to reach equal levels of denoising, one should take a larger regularizing factor λ in the weighted case. We then raise the question: which λ should we take to reach the denoising level of the denoising with unweighted windows and a parameter λ' ? Next proposition states an asymptotic normalization result of λ with respect to the window's smoothness, in a special case. In particular, we need to assume that the domain is discrete and that the total variation is regularized, to be differentiable.

Proposition 5.5 (λ normalizing for smooth windows) *Let (ω_x) denote positive weights associated to the finite window \mathcal{W} , ω_0 being the central weight, and let $T_{\lambda,(\omega_x)}$ be the local TV-MAP denoiser with parameter λ and weights (ω_x) . Assume that the Total Variation is regularized into*

$$TV(u) = \sum_{x \in \mathcal{W}} \sqrt{\varepsilon^2 + |\nabla u(x)|^2}$$

Let $T_{\lambda,(\omega_x)}$ denote the local TV-MAP denoising operator associated to the regularizing parameter λ and the weights (ω_x) . When λ goes to 0, then

$$T_{\lambda,(\omega_x)} = T_{\lambda',(1)} + o(\lambda) \quad \text{where} \quad \lambda' = \frac{\lambda \sum_x \omega_x}{\omega_0 |\Omega|}.$$

This means that when the central weight ω_0 is the maximum weight, and when $\lambda \rightarrow 0$, the denoising operator $T_{\lambda,(\omega_x)}$ is equivalent to the local TV-MAP denoiser with constant weights (non-smooth window), associated to the regularization parameter λ' . Notice that in the Gaussian case, ω_0 is the maximum weight, and hence $\lambda' < \lambda$. We find back that the smooth windows imply lesser denoising.

Proof : Assume that the weights are normalized such that

$$\sum_{x \in \mathcal{W}} \omega_x = |\Omega|.$$

Let $v \in \mathbb{R}^\Omega$. Then $T_{\lambda,(\omega_x)}(v)(x)$ is the central gray level of the $u \in \mathbb{R}^{\mathcal{W}_x}$ which minimizes

$$\sum_{y \in \mathcal{W}_x} \omega_y (u(y) - v(y))^2 + \lambda TV(u), \quad (5.15)$$

and hence satisfies

$$\forall y \in \mathcal{W}_x, \quad \omega_y (u(y) - v(y)) - \frac{\lambda}{2} \operatorname{div} \frac{\nabla u}{\sqrt{\varepsilon^2 + |\nabla u|^2}}(y) = 0,$$

which writes

$$(Id - \frac{\lambda}{2} D^{-1} J) u = v, \quad (5.16)$$

where D is the diagonal operator defined on $\mathbb{R}^{\mathcal{W}_x}$ by

$$\forall y \in \mathcal{W}_x, \quad Du(y) = \omega_y u(y),$$

and J is the linear operator defined on the finite dimensional space $\mathbb{R}^{\mathcal{W}_x}$ such that

$$\forall u \in \mathbb{R}^{\mathcal{W}_x}, \quad Ju = \operatorname{div} \frac{\nabla u}{\sqrt{\varepsilon^2 + |\nabla u|^2}}.$$

The operator $Id - \frac{\lambda}{2} D^{-1} J$ is invertible because the variational formulation (5.15) has a unique solution. Notice that the solution u of (5.16) depends on λ . For any $\lambda > 0$, let us denote

$$w_\lambda = \frac{u - v}{\lambda}.$$

As $u \rightarrow v$ when $\lambda \rightarrow 0$, we get the convergence $\lambda w_\lambda \rightarrow 0$ when $\lambda \rightarrow 0$. Besides, replacing u by $v + \lambda w_\lambda$ in Equation (5.16) yields

$$v + \lambda w_\lambda - \frac{\lambda}{2} D^{-1} J(v + \lambda w_\lambda) = v$$

i.e.

$$w_\lambda = \frac{1}{2}D^{-1}J(v + \lambda w_\lambda).$$

As $\lambda w_\lambda \rightarrow 0$ when $\lambda \rightarrow 0$, and as $D^{-1}J$ is continuous at point v , we get

$$w_\lambda = \frac{u - v}{\lambda} \xrightarrow{\lambda \rightarrow 0} \frac{1}{2}D^{-1}J(v).$$

Fetching the central point of the patch yields

$$u(x) = v(x) + \frac{\lambda}{2}\omega_0^{-1}J(v)(x) + o(\lambda).$$

Denoting $\lambda' = \omega_0^{-1}\lambda$, we obtain

$$T_{\lambda,(\omega_x)}(v)(x) = T_{\lambda',(1)}(v)(x) + o(\lambda).$$

because when the (ω_x) are all equal and normalized, they all equal 1, and $\omega_0^{-1} = 1$. This concludes the proof. \square

5.3.4 PDE associated to local TV-MAP

In this subsection we assume that the images are defined on a continuous domain. The Total Variation energy is regularized and minimized on each disc $B(x, h)$, where $x \in \Omega$ and $h > 0$. The induced denoiser can be considered as a neighborhood filter. Next proposition is about the behavior of the solution of TV-MAP when the radius h of the window goes to zero (with a regularity parameter λ constant, independent on h). We show that the partial differential equation (PDE) associated to this neighborhood filter is the heat equation.

Proposition 5.6 (PDE associated to local TV-MAP) *Let Ω be an open and connected subset of \mathbb{R}^2 . Let $x \in \Omega$ and $h > 0$. Let us assume also that the total variation on $\mathbb{R}^{B(x,h)}$ is regularized and defined by*

$$TV_\varepsilon(u) = \int_{B(x,h) \cap \Omega} \sqrt{\varepsilon^2 + |\nabla u(y)|^2} dy.$$

Let $v \in \mathbb{R}^\Omega$ be a \mathcal{C}^2 image defined on Ω , and let us consider its denoising thanks to the local TV-MAP defined on disks with radius h , denoted $T_h(v)$, i.e. such that $T_h(v)(x)$ minimizes

$$\mathcal{E}_{x,h}(u) = \int_{B(x,h)} (u(y) - v(y))^2 dy + \lambda \int_{B(x,h)} \sqrt{\varepsilon^2 + |\nabla u(y)|^2} dy.$$

Let us also assume that there exists $h_0 > 0$ such that $B(x, h_0) \subset \Omega$ and

$$\sup_{0 < h < h_0} \sup_{y \in B(x,h)} \|\nabla^5 T_h(v)(y)\| < +\infty. \quad (5.17)$$

Then, when h goes to 0, the denoised image $T_h(v)$ satisfies

$$T_h(v)(x) = v(x) + \frac{h^2}{6}\Delta v(x) + o(h^2),$$

where Δv denotes the Laplacian of v .

The proposition implies that the PDE associated to iterations of the local TV-MAP when the window shrinks to a point is the heat equation

$$\frac{\partial u}{\partial t} = \Delta u.$$

Iterating n times the local TV-MAP with $\frac{nh^2}{6} = t$ is equivalent to applying the heat equation during a time t to the input image v , when $h \rightarrow 0$ and $n \rightarrow \infty$.

The assumption (5.17) about the uniform boundedness of $\nabla^5 T_h(v)$ will be discussed after the proof.

Proof : With no loss of generality we can assume $x = 0$. Let then $h > 0$ such that $B(0, h) \subset \Omega$. The solution of the local TV-MAP denoising at 0 comes from the image $u \in \mathbb{R}^{B(0, h)}$ which minimizes

$$\mathcal{E}_h(u) = \int_{B(0, h)} (u(x) - v(x))^2 dx + \lambda \int_{B(0, h)} \sqrt{\varepsilon^2 + |\nabla u(x)|^2} dx.$$

The differential of this energy writes thanks to Stokes' theorem

$$\begin{aligned} d\mathcal{E}(u)(w) = & 2 \int_{B(0, h)} w(x)(u(x) - v(x)) dx - \lambda \int_{B(0, h)} w \operatorname{div} \frac{\nabla u}{\sqrt{\varepsilon^2 + |\nabla u|^2}} dx \\ & + \lambda \oint_{S(0, h)} w \left(\frac{\nabla u}{\sqrt{\varepsilon^2 + |\nabla u|^2}} \cdot \vec{n} \right) dx, \end{aligned}$$

where ∇u and $\operatorname{div}(\nabla u / \sqrt{\varepsilon^2 + |\nabla u|^2})$ are well defined when $h < h_0$, because of Assumption (5.17). The image u which minimizes \mathcal{E} satisfies the Euler-Lagrange equation attached to \mathcal{E} , which is

$$u(x) - v(x) = \frac{\lambda}{2} \operatorname{div} \frac{\nabla u}{\sqrt{\varepsilon^2 + |\nabla u|^2}}(x) \quad (5.18)$$

with the boundary conditions

$$\nabla u(x) \cdot \vec{n}(x) = 0 \quad \forall x \in S(0, h). \quad (5.19)$$

($S(0, h)$ denotes the centered circle with radius h). Attention must be paid to the fact that the solution u to the minimizing problem depends on h . A correct notation for the solution would be u_h but we omit the index h to make them lighter.

Let $x \in B(0, h)$. Then x can be written $x = h\vec{n}$ with $\vec{n} \in S(0, 1)$. A Taylor expansion of $\nabla u(x)$ writes

$$\nabla u(x) \cdot \vec{n} = \nabla u(0) \cdot \vec{n} + h \nabla^2 u(0)(\vec{n})^2 + \frac{h^2}{2} \nabla^3 u(0)(\vec{n})^3 + \frac{h^3}{6} \nabla^4 u(0)(\vec{n})^4 + h^4 \int_0^1 \frac{(1-t)^4}{6} \nabla^5 u(tx)(\vec{n})^5 dt.$$

Thanks to Assumption (5.17), the integral in the equation can be uniformly bounded in $h \in (0, h_0)$. A similar expansion applied on $\nabla u(-x)$ yields, thanks to the boundary condition (5.19),

$$\begin{cases} 0 = \nabla u(x) \cdot \vec{n} = \nabla u(0) \cdot \vec{n} + h \nabla^2 u(0)(\vec{n})^2 + \frac{h^2}{2} \nabla^3 u(0)(\vec{n})^3 + \frac{h^3}{6} \nabla^4 u(0)(\vec{n})^4 + O(h^4), \\ 0 = \nabla u(-x) \cdot \vec{n} = \nabla u(0) \cdot \vec{n} - h \nabla^2 u(0)(\vec{n})^2 + \frac{h^2}{2} \nabla^3 u(0)(\vec{n})^3 - \frac{h^3}{6} \nabla^4 u(0)(\vec{n})^4 + O(h^4), \end{cases}$$

where $O(h^4)$ means a domination by Ch^4 , with C independent from h , which holds for $h \in (0, h_0)$. First summing the two equations and then subtracting them we obtain

$$\begin{cases} \nabla u(0) \cdot \vec{n} + \frac{h^2}{2} \nabla^3 u(0)(\vec{n})^3 = O(h^4), \\ \nabla^2 u(0)(\vec{n})^2 + \frac{h^2}{6} \nabla^4 u(0)(\vec{n})^4 = O(h^3), \end{cases} \quad (5.20)$$

which holds for any \vec{n} in the circle $S(0, 1)$. In particular, as (5.17) implies that $\nabla^3 u(0)$ and $\nabla^4 u(0)$ are bounded on $S(0, h)$ uniformly in h , we have

$$\nabla u(0) = O(h^2) \quad \text{and} \quad \nabla^2 u(0)(\vec{n})^2 = O(h^2) \quad \forall \vec{n} \in S(0, 1). \quad (5.21)$$

Furthermore, applying (5.20) to $\vec{n} = (1, 0)$ and $\vec{n} = (0, 1)$ and summing up yields

$$u''_{xx}(0) + u''_{yy}(0) + \frac{h^2}{6}(u''''_{xxxx}(0) + u''''_{yyyy}(0)) = O(h^3), \quad (5.22)$$

which comes as a link between the second and the fourth derivatives of u . Now coming back to the Euler-Lagrange equation (5.18) and applying it at $x = 0$, we obtain

$$\begin{aligned} u(0) - v(0) &= \frac{\lambda \varepsilon^2 (u''_{xx} + u''_{yy}) + u'^2_x u''_{yy} + u'^2_y u''_{xx}}{2(\varepsilon^2 + u'^2_x + u'^2_y)^{3/2}}(0) \\ &= \frac{\lambda}{2\varepsilon} (u''_{xx} + u''_{yy})(0) + o(h^2) \end{aligned} \quad (5.23)$$

thanks to (5.21). To be able to compute $u''_{xx}(0)$ and $u''_{yy}(0)$ with respect to v in (5.23), let us take the derivatives of the Euler-Lagrange equation (5.18). For instance, differentiating with respect to the first variable leads to

$$u'_x - v'_x = \frac{\lambda}{2} \left[-\frac{3u'_x u''_{xx} (\varepsilon^2 (u''_{xx} + u''_{yy}) + u'^2_x u''_{yy} + u'^2_y u''_{xx})}{(\varepsilon^2 + u'^2_x + u'^2_y)^{5/2}} + \frac{\varepsilon^2 u'''_{xxx} + 2u'_x u''_{xx} u''_{yy} + u'^2_y u'''_{xxx}}{(\varepsilon^2 + u'^2_x + u'^2_y)^{3/2}} \right]$$

(the equality holds on $B(0, h)$), which, again differentiated with respect to the first variable, yields, when taking the value at 0 and taking $h \in (0, h_0)$,

$$-v''_{xx}(0) = \frac{\lambda}{2\varepsilon} u''''_{xxxx}(0) + o(h).$$

Similarly,

$$-v''_{yy}(0) = \frac{\lambda}{2\varepsilon} u''''_{yyyy}(0) + o(h).$$

The two last equations put together yield

$$(u''''_{xxxx} + u''''_{yyyy})(0) = -\frac{2\varepsilon}{\lambda} (v''_{xx} + v''_{yy})(0) + o(h). \quad (5.24)$$

Now combining (5.22) to (5.24) gives

$$(u''_{xx} + u''_{yy})(0) = \frac{h^2}{6} \frac{2\varepsilon}{\lambda} (v''_{xx} + v''_{yy})(0) + O(h^3).$$

This approximation of the Laplacian of u can be used in (5.23), and we obtain

$$Tv(0) = u(0) = v(0) + \frac{h^2}{6}\Delta v(0) + o(h^2).$$

This equation is also true for pixel x instead of 0, and the proposition is proven. \square

Remark 1 The final PDE does not depend on ε , so one could believe that it still holds for the non-regularized Total Variation (where $\varepsilon = 0$). However we have got no evidence for such a result, and the regularizing parameter ε is crucial in the proof. It is likely that the total variation being smooth at 0 is responsible for this blurring behavior. Furthermore, it is noticeable that in the case where $\varepsilon > 0$ the first term in the Taylor expansion does not depend on λ either. Most probably the asymptotic behavior should depend on λ when $\varepsilon = 0$.

Remark 2 The condition (5.17) about the uniform boundedness of $\nabla^5 T_h(v)(y)$ is quite unattractive but is necessary in our proof for technical reasons. Now literature on PDEs [Caffarelli and Cabré 1995] shows that under several conditions, the viscosity solution u of an elliptic PDE very often gains two orders of regularity with respect to the datum v . In particular, if v is $\mathcal{C}^{3,\alpha}$ then it is likely that $T_h(v)$ is $\mathcal{C}^{5,\alpha}$ for a certain α (Hölder regularity). The uniform boundedness of the derivatives seems to be more difficult to obtain in full generality.

Proposition 5.6 tells that for each $t > 0$,

$$T_h^n(v) \xrightarrow[\substack{n \rightarrow \infty \\ nh^2/6=t}]{\quad} G_t * v$$

where T_h^n denotes the local TV-MAP computed on discs with radius h iterated n times, and G_t is a Gaussian kernel with bandwidth t , which is equivalent to applying the heat equation during a time interval t . Figure 5.4 shows experiments on the iterated local TV-MAP denoiser. Each row of the figure corresponds to a certain time interval t in the heat equation, and going to the right means increasing n and decreasing h , with $\frac{nh^2}{6} \approx t$. As in the experiments the domain is discrete, the radius h cannot take any values, that is why the equality $\frac{nh^2}{6} = t$ can only be approximated. Going to the right makes the image look blurry, as if it were convolved by a Gaussian kernel. This means that the limiting PDE most probably remains the heat equation when $\varepsilon = 0$.



Figure 5.4: Iterated local TV-MAP. The window with radius h and the number n of iterations are such that $nh^2/6$ is approximately constant on each row. Going to the right means increasing values of n and decreasing values of h .

5.3.5 Local versus global TV-MAP

Here we aim at comparing locally denoised images for different window sizes. In particular the regularity parameter λ may depend on the window, and should be carefully chained.

λ **chaining** The most natural way of normalizing λ through the different window sizes is to come back to the Bayesian interpretation of the TV minimization.

If the prior on the whole images is $\frac{1}{Z} \exp \left[-\beta \frac{\sum_{x \in \Omega} |\nabla u(x)|}{|\Omega|} \right]$, then the prior on the small images restricted to \mathcal{W} is given by

$$\begin{aligned} p(u(\mathcal{W})) &= \frac{1}{Z} \int_{u(\Omega \setminus \mathcal{W})} \exp \left[-\frac{\beta}{|\Omega|} \left(\sum_{x \in \mathcal{W}} |\nabla u(x)| + \sum_{x \notin \mathcal{W}} |\nabla u(x)| \right) \right] du(\Omega \setminus \mathcal{W}) \\ &= \frac{1}{Z'} \exp \left[-\frac{\beta}{|\Omega|} \sum_{x \in \mathcal{W}} |\nabla u(x)| \right]. \end{aligned}$$

Likewise, if the noise is Gaussian with variance σ^2 on every pixel, the distribution on the noise restricted to a window \mathcal{W} is

$$\begin{aligned} p(v(\mathcal{W})|u(\mathcal{W})) &= \frac{1}{Z} \int_{u(\Omega \setminus \mathcal{W})} \exp \left[-\frac{\sum_{x \in \mathcal{W}} (u-v)^2(x) + \sum_{x \notin \mathcal{W}} (u-v)^2(x)}{2|\Omega|\sigma^2} \right] du(\Omega \setminus \mathcal{W}) \\ &= \frac{1}{Z'} \exp \left[-\frac{\sum_{x \in \mathcal{W}} (u-v)^2(x)}{2|\Omega|\sigma^2} \right] \end{aligned}$$

Hence the MAP estimator deduced from this model is the minimizer of

$$\|u(\mathcal{W}) - v(\mathcal{W})\|^2 + \lambda TV(u(\mathcal{W})) \quad \text{with} \quad \lambda = 2\sigma^2\beta,$$

and λ does not depend on \mathcal{W} . Hence, provided that the L^2 -norm and the total variation are well normalized, λ only needs to be taken constant. This allows extensive comparisons between local and global denoising.

Figure 5.5 shows the denoising of an image for a sequence of windows with growing sizes, and constant λ . For $\mathcal{W} = 15 \times 15$ the obtained image is almost equal to the globally denoised one.

The size of the window as a regularity parameter It is noticeable that the images denoised with a small window contain more noise than the images denoised with a large window (see Figure 5.5).

The size of the window then acts like a regularizing parameter. When the window is very small, the local TV-MAP denoising is practically non effective. Indeed for the extreme case where the window contains one pixel only (i.e. $\mathcal{W}_x = \{x\}$), then the local TV-MAP denoiser is the identity, and does not modify the input image. When the window is very large, with a constant parameter λ , the local TV-MAP denoiser acts more and more like the global TV-MAP denoiser, as shows Figure 5.7.




Figure 5.5: From local to global denoising. Here is shown a noisy image (top left) together with its denoised versions using the local TV-MAP local with different window sizes (hard windows) (with $\lambda = 20$ everywhere). The chain is ended up with the classical global TV-MAP denoising. We can see that the regularity of the denoised image increases when the window size $|\mathcal{W}|$ increases. Furthermore the images denoised with windows as small as 9×9 or 15×15 are very similar to the globally denoised image. It shows that in the denoising of a certain pixel the impact of far away pixels is limited.

This interpretation of $|\mathcal{W}|$ as a regularity parameter can be partially and heuristically explained as follows. Let $v \in \mathbb{R}^{\mathcal{W}}$ denote a noisy patch. The denoised patch is denoted Tv and satisfies

$$Tv = \arg \min_{u \in \mathbb{R}^{\mathcal{W}}} \sum_{x \in \mathcal{W}} (u_x - v_x)^2 + \lambda \sum_{x \in \mathcal{W}} \sum_{\substack{y \sim x \\ y \in \mathcal{W}}} |u_x - u_y|,$$

(with $x \sim y \Leftrightarrow |x - y| = 1$). In particular, when subdifferentiating the energy, we get

$$\forall x \in \mathcal{W}, (Tv)_x \in v_x - \sum_{\substack{y \sim x \\ y \in \mathcal{W}}} \text{sign}((Tv)_x - (Tv)_y),$$

where the set-valued function sign is defined by

$$\text{sign}(y) = \begin{cases} \{1\} & \text{if } y > 0, \\ \{-1\} & \text{if } y < 0, \\ [-1, 1] & \text{if } y = 0. \end{cases}$$




Figure 5.6: From local to global denoising with constant denoising level. The top-left image is denoised using the local TV-MAP local with different window sizes (hard windows). The regularity parameter λ was tuned in each case to reach a desired denoising level (L^2 -distance between the denoised and the noisy image) ($\lambda = 22.47$ for $\mathcal{W} = 3 \times 3$, $\lambda = 19.93$ for $\mathcal{W} = 5 \times 5$, $\lambda = 19.92$ for $\mathcal{W} = 9 \times 9$, $\lambda = 19.97$ for $\mathcal{W} = 15 \times 15$, and $\lambda = 20.00$ for $\mathcal{W} = \Omega$, yielding a denoising level ≈ 9.32 , whereas the noise's s.d. = 10).

In particular,

$$\forall x \in \mathcal{W} \setminus \{0\}, (Tv)_x \in \begin{cases} [v_x - 2\lambda, v_x + 2\lambda] & \text{if } x \text{ is in the interior of } \mathcal{W}, \\ [v_x - \frac{3}{2}\lambda, v_x + \frac{3}{2}\lambda] & \text{if } x \text{ is on a border of } \mathcal{W} \text{ (but not a corner),} \\ [v_x - \lambda, v_x + \lambda] & \text{if } x \text{ is on a corner of } \mathcal{W}, \end{cases}$$

while the central gray level $(Tv)_0$ satisfies $Tv \in v - \frac{\lambda}{2} \sum_{x \sim 0} \text{sign}((Tv)_0 - (Tv)_x)$. When the window \mathcal{W} is large, then the interior of \mathcal{W} contains lots of pixels and the corresponding gray levels $(Tv)_x$ have a broader constraint interval $[v_x - 2\lambda, v_x + 2\lambda]$. Then the neighbors $(Tv)_x$ may become equal more easily, and propagating towards the center, this allows more $(Tv)_x$ with $x \sim 0$ to be equal. Hence the constraint interval for Tv_0 is more likely to be large, and the difference $|v_0 - (Tv)_0|$ is more likely to be high too. The denoising can therefore be more effective, even for a fixed value of λ . Of course there should be better theoretical explanations for \mathcal{W} to be a regularizing parameter.

Now as we saw in the first part of this thesis, the global TV-MAP has the characteristics of a denoiser which denoises too much, as the staircasing shows. For instance, a barycenter between the global TV-MAP and the noisy input generally behaves better than TV-MAP alone, both




Figure 5.7: From local to global denoising. For each of the 4 classical images Cameraman, Lena, Barbara and Goldhill, degraded by Gaussian white noises, the graph shows the decrease of the L^2 -distance between the globally denoised image and the locally denoised one using increasing window sizes (hard window, and constant $\lambda = 20$). Despite the fact that the images are very different (miscellaneous textures with different scales), the decreasing rate seems to be independent of the image. Besides, as this distance converges to zero quite fast, it is clear that even visually the locally denoised images are very similar to the globally denoised image as soon as the window size is large enough, say larger than 13×13 (the plotted values are the average of L^2 -distance values over a sample of 5 noises).

visually and with SNR quality criteria. TV-LSE, the denoising method developed in the first part achieves a better compromise between the TV-MAP and the noisy input than the simple barycenter. Now the localization of TV-MAP is another way to achieve such a compromise. This is why medium-sized windows may be valuable in image denoising. The next subsection deals with the choice of the best window size.

5.3.6 Bias-variance trade-off for window optimizing

Bias-variance decomposition Here we consider the quality of denoising through the expected Euclidean distance between the denoised image and the original image, i.e. the mean squared error. We aim at finding the window size which minimizes this error.

Let u be the original image, and ε a noise. Then $v = u + \varepsilon$ is denoised into $T_{\mathcal{W}}v$, the locally denoised image with window \mathcal{W} . The mean squared error, denoted $e(u, \mathcal{W})$, can be classically

split into a bias and a variance term through

$$e(u, \mathcal{W}) = \mathbb{E}_\varepsilon \|u - T_{\mathcal{W}}v\|^2 = \|u - \mathbb{E}_\varepsilon T_{\mathcal{W}}v\|^2 + \mathbb{E}_\varepsilon \|T_{\mathcal{W}}v - \mathbb{E}_\varepsilon T_{\mathcal{W}}v\|^2.$$

The bias term $\|u - \mathbb{E}_\varepsilon T_{\mathcal{W}}v\|^2$ is increasing when the size of the window increases. Indeed, when the window \mathcal{W} is small, as seen in the previous subsection, $T_{\mathcal{W}}v$ is close to the noisy image v , whose expectation $\mathbb{E}_\varepsilon T_{\mathcal{W}}v$ equals u . This makes a small bias. Its value is 0 when the window only contains one pixel (but the denoising is ineffective). Conversely the bias is large when the window \mathcal{W} is large; more precisely it converges to the bias obtained with the global TV-MAP denoiser, which is positive for natural images. We expect it to be increasing with a slope governed by the proximity of the original image u to the TV model.

The variance term $\mathbb{E}_\varepsilon \|T_{\mathcal{W}}v - \mathbb{E}_\varepsilon T_{\mathcal{W}}v\|^2$ has an opposite behavior: when \mathcal{W} is small, $T_{\mathcal{W}}v(x)$ depends on the few pixels of $v(x + \mathcal{W})$, and then the variance is large. Conversely, when \mathcal{W} is large, many pixels contribute to the denoising of $v(x)$, what reduces the impact of the noise. The variance is then small. We expect it to be decreasing like $1/\sqrt{|\mathcal{W}|}$ as in the central limit theorem.




Figure 5.8: Bias-variance decomposition of the risk, with respect to the window size. The bias and the variance were computed on a sample of 50 noises added to a 100×100 crop of Lena.

This is consistent with the classical bias-variance decompositions which occur in estimation or classification: $|\mathcal{W}|$ indicates a measure of the model’s complexity, and the bias increases and the variance decreases when this complexity gets larger. For every image there must exist a window size which minimizes the mean squared error. We refer at it as the “optimal window” \mathcal{W} .

Optimal window and characteristic scale We now conduct experiments showing that the optimal window size is linked to a “characteristic scale” of the original image u . We mean by a characteristic scale a scale such that in scale-space representation, at a finer scale the structure interferes with spurious texture or noise, and at a coarser scale the salient structure is blurred and vanishes (see [Luo et al. 2007] and references therein). This is a perceptual definition, and in practice the characteristic scale of a natural image is most of the time quite difficult to decide.

This is why we start with synthetic images, defined by

$$u_\omega(x) = 128 (1 + \sin(\omega x_1) \sin(\omega x_2))$$

for several values of the frequency ω (Figure 5.9), brought back to the interval $[0, 255]$. The scale of u_ω is a decreasing function of ω . For each value of ω , we draw a set of 20 noisy versions (v_ω^i) of u_ω , and we estimate the expected squared error by the average of the error on these 20 images

$$e(u, \mathcal{W}) = \mathbb{E}_\varepsilon \|u_\omega - T_{\mathcal{W}}(u_\omega + \varepsilon)\|^2 \approx \frac{1}{20} \sum_{i=1}^{20} \|u_\omega - T_{\mathcal{W}}v_\omega^i\|^2.$$

The values are plotted for several windows \mathcal{W} on Figure 5.10. It comes that the window size which minimizes the error is a decreasing function of the frequency ω . In other words, the greater the scale, the larger the optimal window, which sounds quite natural.

More precisely, the optimal window size $|\mathcal{W}|$ seems to be linearly linked to $1/\omega$, as shows Figure 5.11. It indicates that the optimal window size can be predicted from the image.




Figure 5.9: Synthetic images assumed to have homogeneous scale. For different values of frequency ω , the gray level of a pixel x is given by $u_\omega(x) = 128 (1 + \sin(\omega x_1) \sin(\omega x_2))$ (the images u_ω have normalized gray levels in $[0, 256]$, and contain 128×128 pixels).

Switching to natural images is difficult because very often the characteristic scale is highly location-dependent. For instance an image where fine texture and coarser texture are side to




Figure 5.10: The synthetic images of Figure 5.9 are added with Gaussian white noise and denoised via local TV-MAP denoiser with different window sizes (constant $\lambda = 20$). The graph represents the L^2 -distance between the locally denoised images and the original image u_ω along the different window sizes. For each ω the curve reaches a minimum, defining the optimal window size. This optimal window size is intimately correlated to the frequency ω . In particular it decreases when ω increases (the values on the graph are averages over a sample of 20 experiments).

side intuitively has two main scales: the small scale at the location of the fine texture and the larger scale at the location of the coarser texture. For the moment we consider pieces of images for which the scale is more or less homogeneous, and give experimental evidence for the fact that the optimal window is an increasing function of the scale. The case of heterogeneous scale will be tackled in the next subsection.

First a smooth zone of a natural image was cropped (corresponding to very large scale); then a region composed of quite large smooth zones separated by contrasted edges; then two textures with two different small scales. Noisy versions of these images are shown on Figure 5.12 (top row), together with their denoised versions for different windows \mathcal{W} . The best visual quality is obtained for larger \mathcal{W} for the smooth image, while smaller \mathcal{W} become better for thin texture. In other words, the window achieving the best visual quality is smaller when the scale is finer. This also holds for the mean squared error $e(u, \mathcal{W})$: the window minimizing $e(u, \mathcal{W})$ is small when the scale is small and vice versa (Figure 5.13) ($e(u, \mathcal{W})$ was approximated by the average over a sample of 20 noisy versions of u).

Assuming that the scale is proportional to the optimal $|\mathcal{W}|$, we obtain for the smooth image a scale proportional to 441 for the smooth image, 25 for the piecewise constant image, 9 for the coarse texture and 1 for the fine texture. These values counted in pixels are a good intuitive measure of the scale of the related images.

For arbitrary images containing several characteristic scales, we can expect that the choice




Figure 5.11: Fit between the optimal window size $|\Omega|$ obtained in the denoising of the synthetic images u_ω shown on Figure 5.9 and the frequency inverse $1/\omega$. The inverse of ω represents the scale of the original u_ω , and the graph shows that the optimal window size (measured by the number of pixels it contains) is approximately a linear function of this scale. The fit error is very large in the small scales because in this case the pixel size requires a very rough quantization of the window size.

of a window corresponding to one of these scales will properly denoise certain parts of the image and others will be really damaged. This is why a locally defined scale gets necessary. From these local scales a location-dependent optimal window can be derived, and a new denoiser is defined by denoising each pixel thanks to the optimal window attached to this pixel. In the next subsection we give details about the estimation of the local scale, and derive window-adaptive local TV denoising.




Figure 5.12: Local TV-MAP denoising with different window sizes on natural images. The selected images are such that the scale is approximately homogeneous. The top row shows noisy versions of them (60×60 crops of Lena, Peppers, again Lena and Barbara, and noise's s.d. $\sigma = 10$), sorted by decreasing “intuitive characteristic scale”. Each following row shows the local denoising of the noisy image using a particular window size (3×3 for the 2nd row, 7×7 for the third row, and 21×21 for the last one, with always $\lambda = 20$). Considering each column separately, which image is visually the best? For the first column, as the original is perfectly smooth, the greatest level of denoising is preferable, i.e. a large window; however the staircasing effect even in the local denoising makes it difficult to maintain the largest window as the optimal one. We would rather choose the 7×7 window. In the second column, the objects are big and well contrasted, and again the 7×7 window seems best. The texture in the third column is too much erased by large windows, but the noise is clearly visible in the noisy image; the best window is probably the 3×3 . Last but not least, the thin texture from the 4th column is severely damaged in any denoised version, and contrast is shrunk; no doubt the best image is the noisy one, corresponding to a 1×1 optimal window.

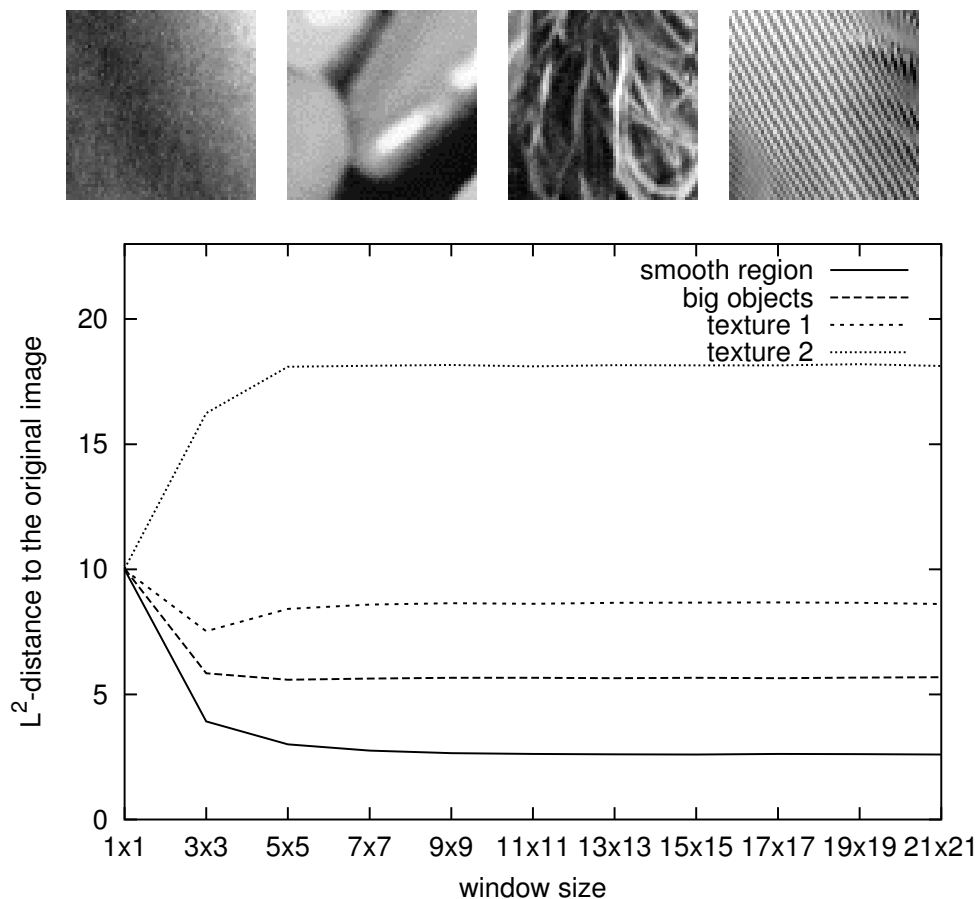


Figure 5.13: The 4 original images shown on the top row are noised ($\sigma = 10$) and then denoised via the local TV-MAP denoiser for several window sizes ($\lambda = 20$). The graph shows the L^2 -error between the locally denoised image and the original against the window size. The optimal window related to the smooth image is very large; it is smaller for the second image made of large objects. The third image representing a texture has a 3×3 optimal window, while the thin stripes of the 4th image have a 1×1 optimal window, which means that the best is to operate no denoising at all). This graph comes as a confirmation of the intuitive optimal windows commented on Figure 5.12, except for the smooth image. This exception can be understood by the fact that the L^2 -error, as well as SNR quantities, do not take visual artefacts such as staircasing effect into account. Furthermore it seems again that the “scale” of the images is negatively correlated with the optimal window size.

5.3.7 Adaptive locality as a function of the gradient norm

Introduction

In this subsection we aim at finding a practical estimator for the window size which achieves the best denoising, and to use it in the local TV-MAP denoiser. This optimal window size needs to be locally adapted because usually the images contain objects with different sizes, and structures with different characteristic scales, which determine different sizes for the optimal windows. This means that the corresponding local TV-MAP denoiser needs to be necessarily locally adaptive.

This idea of locally adaptive TV-based denoising has already been explored by several authors, namely [Almansa et al. 2008; Bertalmio et al. 2003; Gilboa et al. 2003; Moisan 2001; Strong et al. 1997], particularly through the local adaptation of the regularity parameter λ . For instance in [Strong et al. 1997], the global energy

$$\|u - v\|^2 + \sum_x \lambda(x) |\nabla u(x)|$$

is minimized, where λ depends on x and is chosen as

$$\lambda(x) \propto \frac{1}{|\nabla u(x)|} \quad \text{where } u \text{ is a rough estimate of the original image.}$$

This choice for $\lambda(x)$ is such that it will be small when the image oscillates much and then it will operate little denoising, while it will be large when u is smooth, which will allow maximum denoising.

In [Almansa et al. 2008] a Mumford-Shah segmentation of the noisy image is run. λ is assumed to be constant on each region of the partition, and is tuned in order to get the desired level of denoising.

Here a similar procedure can be done on the window size $|\mathcal{W}|$ also considered as a regularity parameter. As seen in Subsection 5.3.6, the optimal window is deeply linked to the local scale, so the choice for $|\mathcal{W}|$ is a problem of local scale estimation. This subsection is organized as follows: first we briefly comment the local scale estimation, and show that the inverse of the gradient norm, as $\lambda(x)$ in [Strong et al. 1997] can be a very handy (but limited) estimate of it.

Then a space-adaptive local TV-MAP denoiser is derived from the locally optimal windows, and discussed. In particular this denoiser is able to achieve good denoising on heterogeneous but simple images while assuming a constant Bayesian model. However this approach shows severe visual artefacts on natural images.

Estimation of the optimal window by the local scale

As seen in the previous subsection, the window which achieves the best local denoising is a function of the characteristic scale. Then a scale estimator is enough to make the optimal window size available.

The scale of an image can be mathematically defined in several ways. For images drawn from Gibbs fields, the scale can be thought of as the maximal size of the cliques. In a Markov random field interpretation, it is about the same to consider the basic neighborhood of the field. [Csiszar

and Talata 2006] builds an estimator of the size of this basic neighborhood, based on a penalized likelihood maximization. But this estimator does not apply for non-stationary processes such as natural images. And even a scale estimate computed on a small part of an image is expected to be very bad due to the small sample size. Hence an efficient local scale estimate seems difficult to get with such a method.

[Luo et al. 2007] propose a reliable resolution-independent estimate for the characteristic scale. The characteristic scale is estimated by the one which maximizes a normalized total variation over the scale space of the image. Localization is possible. We also mention [Brox and Weickert 2006] where the local scale is measured from the evolution speed of the gray levels submitted to a TV flow.

Now we rather quit the scale interpretation of the optimal window and try to estimate it directly.

Definition of a locally optimal window

Let us try a definition of a locally optimal window. Let then x denote a pixel. The optimal window $\mathcal{W}(x)$ centered at x is requested to reach the smallest possible error in the denoising of pixel x . Its radius $\rho(u, x)$ is defined by

$$\rho(u, x) = \arg \min_{\rho > 0} \mathbb{E}_n |u(x) - T_{\mathcal{W}}(u + n)(x)|,$$

where the expectation is over the noise n . Then $\rho(u, x)$ does not depend on the noise, and achieves the minimal error in average.

The image $x \mapsto \rho(u, x)$ is called here the locality map of u . Figures 5.14 and 5.15 show the locality map of images computed on a sample of one noise, and then on 50 noises (top rows).

Optimal locality and gradient norm

The locality map (Figure 5.15) looks very much like a decreasing function of the gradient norm of the image: when the image is smooth ρ is large, and vice versa. Next proposition proves that the unique function correlating the optimal window size to the gradient norm compatible with scale invariance is the inverse function.

Proposition 5.7 (Continuous framework) *Assume the domain Ω is continuous and equals \mathbb{R}^2 , and let u be an image defined on Ω . Assume there exists a function $f : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ such that $\rho(u, x) = f(|\nabla u(x)|)$ for every u and x . If ρ is scale invariant, i.e.*

$$\forall \lambda > 0, \quad \rho(u(\lambda \cdot), \lambda x) = \lambda \rho(u, x),$$

then there exists $a \in \mathbb{R}$ such that

$$\forall y \in \mathbb{R}_+^*, \quad f(y) = \frac{a}{y}.$$

Proof : Let $u \in \mathbb{R}^\Omega$ and $x \in \Omega$, such that $|\nabla u(x)|$ exists and is positive. For every $\lambda > 0$ let u_λ the image u dilated by a factor λ , i.e.

$$\forall x \in \Omega, \quad u_\lambda(x) = u(\lambda x).$$

Then for all $\lambda > 0$,

$$\frac{\rho(u, x)}{\lambda} = \rho(u_\lambda, \frac{x}{\lambda}) = f\left(|\nabla u_\lambda|(\frac{x}{\lambda})\right) = f(\lambda|\nabla u|(x)).$$

For u and x fixed, the expression $\lambda f(\lambda|\nabla u|(x))$ does not depend on λ . Let $a(u, x) = \frac{1}{|\nabla u|(x)}$. Then $f(y) = a(u, x)/y$ for every $y > 0$. As assumed in the proposition f does not depend on u or x . Then $a(u, x) = a$, and we obtain the desired result. \square

Therefore the simplest approximation of $\rho(u, x)$ from the gradient norm is given by

$$\hat{\rho}(u, x) \propto \frac{1}{|\nabla u(x)|}. \quad (5.25)$$

In Figure 5.14 and 5.15 (bottom row) we compare the locality map to an estimate of the gradient norm's inverse. We choose the following scheme proposed by Moisan [2007]

$$w_\mu(u, x) = \sup \left\{ r \mid \mu \geq \inf_{|y-x| \leq r} \sup_{|z-y| \leq r} |u(z) - u(x)| \right\}$$

which is a non-local estimate of $\mu/|\nabla u(x)|$ when $\nabla u(x) \neq 0$, more convincing than $(\epsilon + |\nabla u(x)|)^{-1}$ or $|\nabla u * G_\sigma(x)|^{-1}$. The main features of the locality map can be retrieved in the gradient norm's inverse. When blurring a little the locality map (which is noisy because the expectation in the locality map was computed on a finite sample) and considering the gradient norm's inverse computed on the original noise-free image, the correlation is slightly larger than 0.85 (Figure 5.14). This correlation decreases a lot (≈ 0.64) when computing the gradient norm's inverse directly on the noisy image.

Notice that a good scale estimate should be contrast invariant, while here the inverse gradient norm is highly contrast dependent. This could be considered as a major drawback of the method. However contrast invariance is not completely true when dealing with optimal window rather than scale, because in a denoising issue the noise's intensity is of course contrast dependent, which compels the true optimal window to depend a little on the contrast.

Now this optimal window estimation by inverse gradient norm contains a severe limitation. Namely in our approach the original image is needed to provide a correct estimate for ρ , because when replacing u with its noisy version really spoils the quality of estimation (see Figures 5.14 and 5.15, bottom rows). We could circumvent this problem by post-processing the locality map, e.g. by convolving the noisy locality map by a Gaussian kernel

$$\hat{\rho}(u, x) = G_g * w_\mu(v, \cdot)(x).$$

for instance. Even sophisticated methods, such as pre-estimations of the original image before computing its inverse gradient norm, could lead to more powerful estimates.

Yet these additional processings may introduce extra parameters. Here we do not aim at providing a complete recipe for adaptive denoising. Therefore we limit ourself to the simplest but unrealistic condition: we assume that the original image u can be used in the locality computation.

Adaptive denoising

Now we propose to adaptively denoise an image v according to the local context. Namely, a pixel x is denoised by

$$T(v)(x) = u(0) \quad \text{where} \quad \begin{cases} u \in \mathbb{R}^{B(0, \hat{\rho}(v, x))} \text{ minimizes } \|u - v(B(x, \hat{\rho}(v, x)))\|^2 + \lambda TV(u), \\ \hat{\rho}(v, x) = w_\mu(u, \cdot)(x). \end{cases}$$

Each pixel is then denoised with its “personal” optimal window. The pixels belonging to high gradient regions, probably corresponding to thin and contrasted textures, will naturally be less denoised than the pixels belonging to a smooth region. The amount of local denoising is controlled by the window size rather than by the regularity parameter λ . This is quite natural, because the information brought by a large homogeneous region is useful for the denoising of each of its pixels, while in more random configurations information provided by far away pixels is virtually useless.

We show experiments conducted on natural images, over non-smooth windows (see Subsection 5.3.2). As the window size needs to be discretized into 1×1 , 3×3 windows and so on, the denoising level is quantified, and the denoising is necessarily very heterogeneous. For example two neighbor pixels, even with close gradient norm values, can be denoised with very different window sizes, and yield highly inhomogeneous denoising. This means that the method noise, i.e. the noise remaining from the difference $v - T(v)$, is far from satisfying the statistical features of a white noise.

Even a regularization of the inverse gradient norm estimate does not manage to circumvent this heterogeneity. Some results are shown on Figure 5.16, and focus on the artefacts of such a method.

To put things in a nutshell, we identify two main artefacts. The first one is a serious problem of denoising heterogeneity. This issue can probably be solved by using a smooth-window version of this adaptive denoising, since then this quantization effect no longer occurs. Another artefact is the leak of thin texture’s average through their boundaries, as in Figure 5.16 (third row). This is explained by the fact that pixels belonging to a smooth part of the image are associated to a large optimal window, even if the pixels lie near the boundary separating the smooth part from a finely textured region. Then the denoising of these pixels depends on pixels belonging to the textured region, and its gray level is necessary attracted by the texture’s average value. This artefact could certainly be avoided by using windows which would be moved away from their center, at the price of an increased complexity in the optimal window model.

In the following section we propose another framework which automatically achieves a kind of adaptive denoising.




Figure 5.14: Locality map of a heterogeneous natural image (I). (a) Input image, taken from Lena. It can be partitioned into mainly two or three regions with different scales. (b) The locality map, i.e. the image for which each gray level corresponds to the size of the optimal window. It was computed from a sample of 50 experiments (50 noises). Bright gray level means large optimal window. A brighter zone appears on the left of the image, and dark pixels located on the stripes (contrast is enhanced and some gray levels are then saturated). (c) Two main regions really appear when the average locality map (b) is a bit smoothed (Gaussian kernel with s.d. $g = 2$): a bright region for large scale and a dark region for small scale. (d) A correct estimate of this locality map can be given by the inverse of gradient norm, here computed on the original image. Correlation to (c) ≈ 0.85 . (e) However using the original image is cheating, and here is shown the inverse of gradient norm computed on the noisy image. The correlation to (c) is ≈ 0.64 . A slight smoothing by a Gaussian convolution makes the correlation increase, and is usable for adaptive denoising.




Figure 5.15: Locality map of a heterogeneous natural image (II). (a) Input image, taken from Barbara. It can be partitioned into two main regions with different scales. (b) The locality map, i.e. the image where gray levels correspond to the size of the optimal window. It was computed out of a sample of size 50 (50 noises). A light zone appears on the north-east of the image, as well as dark pixels located on the stripes (same contrast as in (b)). (c) Convolution of (b) with a Gaussian kernel with s.d. $g = 2$. (d) Inverse of gradient norm computed on the original input. The correlation to (c) is ≈ 0.90 . Notice that the dark side of the stripes are darker than in the locality maps, which means that the gradient norm inverse can sometimes be inconsistent. (e) Inverse of gradient norm computed on the noisy image. The correlation to (c) falls to ≈ 0.84 , but a slight smoothing of it is able to give a correct estimate of the locality map.




Figure 5.16: Adaptive Lagrangian local TV-MAP. On the left column, crops of noisy images are shown. On the central column the inverse gradient norm estimate w_μ is plotted. On the last column, the image obtained by the adaptive denoising of the noisy image (1st column) using a quantified version of the locality map (2nd column) is shown. Many artefacts are visually very annoying. The top image is quite well denoised, and both regions have homogeneous denoising. The result is quite good except from the small region lying on bottom-right, where the texture is abruptly destroyed. On the second row, the lips suffer from the window size quantification, and the denoising is not homogeneous. The third row shows an artefact coming on a contrasted boundary. Both regions are quite well denoised, but the textured part leaks through the contour and makes a whiter zone. This is due to the fact that in the flat region the optimal window size is estimated as very wide, even at the proximity of the boundary. Then pixels near the boundary depend on the other region, and become whiter. The last row is another example of heterogeneous denoising.

5.4 Constrained local TV-MAP

In the previous section the local prior

$$\mathbb{P}(u(\mathcal{W}_x)) = \frac{1}{Z} \exp(-\beta TV(u(\mathcal{W}_x)))$$

was frozen over all the windows of the image. In particular, the parameter β was chosen constant all over the image, even in the adaptive case where the size of the window was changing. Now as the natural images are heterogeneous, the regularity of the image is not uniform, which can be derived into the prior distribution as a TV prior where β is spatially dependent.

It is arguable that the Total Variation is somewhat already multiscale, as shown by the inclusions in the Besov spaces $B_{1,\infty}^1 \subset BV \subset B_{1,1}^1$ [Mallat 1998], and hence should be intrinsically able to deal with multiple regularity. Nevertheless experiments on both TV minimization and TV-LSE show that a small λ in

$$\mathcal{E}(u) = \|u - v\|^2 + \lambda TV(u)$$

is more efficient on thin structures while large λ is better on smooth regions associated with large scales.

More precisely, the selection of a spatially varying λ can be explained by the constrained formulation of the TV minimization. Namely, the Bayesian denoising of a window $v(\mathcal{W}_x)$ must also be consistent with the noise's distribution. In particular, if the noise is white and Gaussian with s.d. σ , the denoised image $T(v)(\mathcal{W}_x)$ is expected to satisfy

$$\|T(v)(\mathcal{W}_x) - v(\mathcal{W}_x)\|^2 \approx \sigma^2$$

provided that the window is large enough to be able to apply the Law of Large Numbers on the variables $(T(v)(x) - v(x))^2$. This can be considered as a local constraint added on the local Total Variation minimization problem, and this constraint results in a Lagrange multiplier $\lambda(x)$ depending on the location x in the optimization problem

$$\text{minimize } \|u(\mathcal{W}_x) - v(\mathcal{W}_x)\|^2 + \lambda(x) TV(u(\mathcal{W}_x)).$$

To sum things up a local TV-MAP denoising associated to the constrained formulation of TV minimization leads to the denoiser T defined by

$$T(v)(x) = w(0) \quad \text{where } w = \arg \min_{w \in \mathbb{R}^{\mathcal{W}}} TV(w) \text{ s.t. } \|w - v(\mathcal{W}_x)\|^2 = \sigma^2. \quad (5.26)$$

when the locality is based on non-smooth windows, and

$$T(v)(x) = w(0) \quad \text{where } w = \arg \min_{w \in \mathbb{R}^{\mathcal{V}}} TV(w) \text{ s.t. } \|w - v(\mathcal{W}_x)\|_{2,\omega}^2 = \sigma^2. \quad (5.27)$$

in the case of smooth windows. This approach is not far from the one of [Almansa et al. 2008] where the total variation of the whole image is minimized under a large set of constraints

$$\|u(x + \cdot) - v(x + \cdot)\|_{2,a}^2 \leq \sigma^2, \quad \forall x \in \Omega. \quad (5.28)$$

The main difference is that our approach is strictly local with a single constraint at a time (and hence much simpler to handle). Our denoiser cannot bring the reliability of a global optimum as in [Almansa et al. 2008]: for example in full generality the local constraints (5.28) may not be satisfied by the denoised image $T(v)$.

5.4.1 Algorithm

The numerical computation of T is again deduced from [Chambolle 2004]. Namely, to solve the problem (5.26) in the case of non-smooth windows, Chambolle's algorithm adapted to the constrained TV minimization can directly be applied on every patch. It simply consists in alternately solving the unconstrained problem with a given λ and then readjusting λ thanks to a linear interpolation between the previous value and the level of denoising $\|u - v(\mathcal{W}_x)\|^2$ which we want to equal σ^2 . The iterations are stopped when the constraint is reached within a precision of 1% or so, and resume the procedure with another pixel. The algorithm can be quite fast when selecting correct initializations:

- First, between 2 iterations of the algorithm over the same window \mathcal{W}_x , the current image u is not reinitialized arbitrarily when coming back to the unconstrained problem solving: the denoised image obtained with the previous λ is a good initializing image, and allows convergence within few iterations.
- Secondly, the pixels are scanned in a lexicographic order, or in any order preserving some kind of continuity in the image. When incrementing the pixel index, the initialization of λ is selected as its last value when the algorithm was computing the previous pixel. As the images have a certain local coherence along their columns (or rows), very often this value is a very good initialization, and a little number of linear interpolations of λ are enough to get a good precision.

Under these conditions, the constrained version of local TV-MAP is almost as fast as the unconstrained version. The procedure is summed up in Algorithm 5.

Algorithm 5 Constrained local TV-MAP algorithm

```

choose an initial precision  $p = 0.01$ ,
choose an initial value for  $\lambda$  ( $\lambda = \sigma^2$ ),
choose an initial patch  $u \in \mathbb{R}^{\mathcal{W}}$  ( $u = 0$ ),
 $r = 1$ ,
for all  $x \in \Omega$  (lexicographical order) do
  repeat
     $\lambda \leftarrow r \cdot \lambda$ 
     $u \leftarrow \arg \min_{u \in \mathbb{R}^{\mathcal{W}}} \|u - v(\mathcal{W}_x)\|^2 + \lambda TV(u)$  (Chambolle's algorithm initialized with  $u$ ),
     $r \leftarrow \frac{\sigma}{\|u - v(\mathcal{W}_x)\|}$ ,
  until  $|r - 1| < p$ .
   $T(v)(x) = u(0)$ 
end for
return  $(T(v)(x))_{x \in \Omega}$ .
```

In the case of smooth windows, Chambolle's algorithm is expected to converge provided that a restriction over the step size is satisfied (see Subsection 5.3.2). Besides, the linear interpolation for λ at each iteration can be proven to yield convergence towards the desired image, because [Chambolle 2004, Theorem 4.2] holds for smooth windows too. Finally Algorithm 5 converges to the required denoised image $T(v)$.

5.4.2 Minimum size of windows?

The window size cannot be arbitrary as in the unconstrained formulation. Indeed, when the window \mathcal{W} contains very few pixels, the approximation

$$\|n(\mathcal{W})\|^2 \approx \sigma^2,$$

n being a Gaussian white noise with s.d. σ , does not hold any more. The same computation as in [Almansa et al. 2008], based on the Central Limit Theorem, shows that the radius of \mathcal{W} should be larger than 10 pixels at least for the constraint to be realistic, when $\sigma \approx 10$. In practice the condition

$$\text{var } v(\mathcal{W}_x) \geq \sigma^2, \quad (5.29)$$

telling that the image is not overwhelmed with noise, which is necessary to the uniqueness of the solution in (5.26) and is also necessary to apply Chambolle's algorithm, does not necessarily hold when the window is too small. The condition 5.29 is even more often violated when the window is small.

In spite of these remarks, in our experiments we investigated windows as small as 3×3 . Indeed the condition (5.29) is not always true, and when it is false for a window \mathcal{W}_x in the image, we chose to assign to $T(v)(x)$ the value

$$T(v)(x) = \frac{1}{|\mathcal{W}|} \sum_{y \in \mathcal{W}_x} v(y).$$

It corresponds to the central gray level of the most realistic solution of the local problem with the inequality constraint

$$\text{minimize } TV(w) \quad \text{s.t.} \quad \|w - v(\mathcal{W}_x)\|^2 \leq \sigma^2,$$

whose solutions w are constant patches, but the most likely is $\frac{1}{|\mathcal{W}|} \sum_{y \in \mathcal{W}_x} v(y)$ because it respects v 's average.

5.4.3 Experiments

The denoising induced by this constrained local TV-MAP algorithm is expected to be somehow adaptive to the local image configurations. For example a thin texture can easily be totally erased by applying TV minimization with a certain λ . The constraint compels the method noise to remain below the threshold σ^2 , and hence enforces the corresponding Lagrange multiplier λ to be small. This can be linked to a property of the Total Variation flow

$$\frac{\partial u}{\partial t} = \text{div} \frac{\nabla u}{|\nabla u|},$$

which is the partial differential equation associated to TV minimization when $\lambda \rightarrow 0$, and shares several properties with TV minimization [Aureu et al. 2000]. [Brox and Weickert 2006] notices that the pixel of an image submitted to the TV flow has a moving speed depending on whether the pixel is extremal or not. In a sophisticated texture many pixels have extremal gray levels and

are then associated to high speed of extinction. This would explain why λ in TV minimization, corresponding to a stopping time in the TV flow, should be small. Conversely pixels lying on smooth zones are associated to low extinction speeds and hence to high values of λ . This property gives some adaptivity to this denoising method.

Figures 5.17 to 5.20 show results of this constrained denoising on classical images, with increasing window sizes. The noise added has s.d. $\sigma = 10$, and this value is kept in the constraint (5.26). The denoising is performed with non-smooth windows. Each denoised image (odd rows) is shown with its corresponding mapping $x \mapsto \lambda(x)$ (even rows). When the condition (5.29) is not satisfied, the associated value of λ is set to the maximum value of λ occurring on the other pixels. For very small window sizes, the image $\lambda(x)$ contains many fluctuations, probably due to the fact that the window is not large enough for the Law of Large Number to apply everywhere. However the regions where λ is high approximately corresponds to smooth zones, while lower values of λ correspond to edges of turbulent zones. When the window size is increased, the mapping $x \mapsto \lambda(x)$ gets smoother, until the size is infinite (classical TV denoising) for which the mapping is constant.

The denoised images are interesting to observe. For small window sizes the denoising is not very homogeneous: this is due to the fact that the mapping λ is very noisy. Therefore nearby pixels are not associated to close values of λ , and spurious noise remains. For large window sizes the denoising is satisfying. Staircasing effect is present, even though each pixel optimizes a different criterion. The locally denoised images slightly outperform the globally denoised one (classical TV minimization, with λ constant). The boundaries and exceptional features are well treated, and the global aspect is satisfying. When the original image is approximately homogeneous, local and global denoising act similarly.

A table of SNR is given on Figure 5.21. It shows that on our examples local denoising outperforms the global TV denoising.

Finally this denoising operator achieves quite good results compared to classical TV denoising. Nevertheless it is not completely satisfactory. For example staircasing effect is still present as though each pixel optimizes a different criterion from its neighbor. The localized denoiser does not prevent either from the isolated pixel artefact which occurs with global TV denoising. Next section concludes by proposing another local and TV-based denoiser which is expected to avoid these artefacts.




Figure 5.17: Constrained local TV-MAP denoising (I). A noisy detail of Barbara (top left) is denoised using the constrained local TV-MAP operator with $\sigma = 10$ and different window sizes. Under each denoised image, the corresponding mapping $x \mapsto \lambda(x)$ used in Chambolle's algorithm is shown (white values mean high values of $\lambda(x)$). On these mappings contrast has been enhanced, giving rise to saturation. Note that when the condition (5.29) is not satisfied, the associated λ was set to $\max \lambda(x)$. For small windows, the condition is very often violated, and this explains why the image of λ is very saturated. Then the denoising lacks homogeneity. When the window size increases, the image of λ gets smoother, and the denoised image is more satisfying. For global denoising of this detail of Barbara image, the λ image is constant ($\lambda \approx 12.8$), and yields a more irregular elbow, with contrasted isolated pixels.




Figure 5.18: Constrained local TV-MAP denoising (II). A noisy detail of House (top left) is denoised using the constrained local TV-MAP operator with $\sigma = 10$ and different window sizes. For each denoised image, the related $x \mapsto \lambda(x)$ is plotted below. In the local denoising, the wall is less denoised than the sky, contrary to the global denoising. This is a good point since the wall contains texture whereas the sky is almost constant, but it leads to contrasted isolated pixels which do not really fit to the original texture.




Figure 5.19: Constrained local TV-MAP denoising (III). A noisy detail of Lena (top left) is denoised using the constrained local TV-MAP operator with $\sigma = 10$ and different window sizes. For each denoised image, the related $x \mapsto \lambda(x)$ is plotted below. For small window sizes, the separation of textures is significant, but it vanishes when the window is too large, yielding virtually constant λ . However for small windows, the textured region contains little artefacts (isolated pixels again). The results are not significantly different from global TV-denoising.




Figure 5.20: Constrained local TV-MAP denoising (IV). A noisy detail of Lena (top left) is denoised using the constrained local TV-MAP operator with $\sigma = 10$ and different window sizes. The eye corresponds to lower values of λ , while the cheek to higher values of λ . An insight comparison between the images shows that staircasing is a bit less marked in locally denoised images than the globally denoised one.

	crop of Fig. 5.17	crop of Fig. 5.18	crop of Fig. 5.19	crop of Fig. 5.20
$\mathcal{W} = 1 \times 1$	13.43	15.74	12.05	11.45
$\mathcal{W} = 3 \times 3$	13.09	20.96	18.48	17.17
$\mathcal{W} = 7 \times 7$	13.58	20.99	18.53	16.89
$\mathcal{W} = 11 \times 11$	13.68	20.72	18.34	16.59
$\mathcal{W} = 15 \times 15$	13.70	20.65	18.27	16.46
$\mathcal{W} = 21 \times 21$	13.71	20.61	18.25	16.38
$\mathcal{W} = \Omega$	13.17	20.73	18.37	16.21

Figure 5.21: SNR values for Figures 5.17 to 5.20. The SNR values were computed on the cropped images only. The window $\mathcal{W} = 1 \times 1$ corresponds to the raw noisy images. For each original image, the “best” denoising is related to a local denoising. It seems then on these examples that constrained local denoising achieves quite good performances, as well visually and on a L^2 criterion.

5.5 Towards a quadratic risk as in the Non-Local means

These TV-based local denoisers can be adapted to a least square error (LSE) criterion, as in the first part of the thesis. In the Lagrangian formulation for instance, an image v is denoised into $T(v)$ satisfying

$$T(v)(x) = \arg \min \mathbb{E}[\|T(v) - u\|^2 | v(\mathcal{W}_x)] \text{ with } \begin{cases} u \sim \frac{1}{Z} e^{-\beta TV(u)} \\ v - u \sim \frac{1}{Z'} e^{-\|v-u\|^2 / (2\sigma^2)}. \end{cases}$$

The conditional expectation assumes that only the neighborhood $v(\mathcal{W}_x)$ of v is known. This is the locality condition. This leads to the explicit formula defining the local TV-LSE denoising operator

$$T(v)(x) = \frac{\int_{\mathbb{R}^{\mathcal{W}}} p(0) e^{-\frac{\|p-v(\mathcal{W}_x)\|^2 + \lambda TV(p)}{2\sigma^2}} dp}{\int_{\mathbb{R}^{\mathcal{W}}} e^{-\frac{\|p-v(\mathcal{W}_x)\|^2 + \lambda TV(p)}{2\sigma^2}} dp},$$

where the local posterior distribution is averaged over the patches $p \in \mathbb{R}^{\mathcal{W}}$. This denoising operator seems quite interesting for two reasons mainly:

- As seen in Chapter 1, the LSE criterion seems pretty much more adapted to image denoising than a MAP criterion. It is likely to avoid troublesome artefacts and statistical distortions.
- The link to the Non-Local means (NL-means) is obvious if T is written

$$T(v)(x) = \frac{1}{Z} \int_{\mathbb{R}^{\mathcal{W}}} P(p) e^{-\frac{\|p-v(\mathcal{W}_x)\|^2}{2\sigma^2}} p(0) dp.$$

Let $P(p) = \frac{1}{Z} e^{-\frac{\lambda}{2\sigma^2} TV(p)}$, i.e. a TV prior on the patches. Then we get the local TV-LSE denoiser. Or else, letting $P(p) = \frac{1}{Z} \sum_{y \in \Omega} \mathbb{1}_{v(\mathcal{W}_y)}(p)$, i.e. associated to an empirical prior on

patches, T exactly becomes the Non-Local means denoising operator. The only difference between this local TV-LSE denoiser and the NL-means lies in the definition of the prior distribution. A comparison would be most valuable.

However to be able to be computed, this denoising operator requires adapted numerical methods. Indeed, the MCMC method described in Chapter 2, which was very well adapted to global denoising, is not necessarily a good choice for local denoising since the approximations we made in high dimension do not hold any more on small patches. In particular, MCMC methods are especially adapted to high dimension, and switching to low dimension means efficiency reduction; the stopping criterion described in Chapter 2 Section 2.3 is probably no longer reliable. A constrained formulation for local TV-LSE raises the same issues.

This point has not been studied in the thesis. Hopefully for very low dimensions (i.e. 3×3 patches or even smaller) direct computations could be achieved, and for intermediate dimensions adapted numerical schemes can be found. It would be interesting then to study the denoising properties of this operator. This could be the object of future work.

Anyway the only difference left between the local TV-LSE denoiser and the NL-means lies in the definition of a prior on patches. This naturally introduces next chapter, where the design of a prior distribution is discussed.

Chapter 6

Total Variation prior model, and the Non-Local means

Résumé Dans ce chapitre nous discutons le choix de la variation totale (TV) comme modèle a priori. Dans un premier temps nous analysons le rapport entre ce modèle et les statistiques locales des images. Nous montrons que la restriction du modèle TV à de petits patches 2×2 est une bonne approximation de la loi de ces patches. En revanche le choix du schéma numérique pour la norme du gradient se révèle assez crucial, et justement les schémas compatibles avec des propriétés géométriques de la variation totale (formule de la coaire, ou bien existence d'une divergence duale du gradient) ne sont pas ceux qui maximisent l'adéquation avec la loi des patches, ce qui fait s'opposer le point de vue géométrique et le point de vue statistique. Puis nous replaçons le modèle TV dans une perspective d'estimation de loi des patches, où les moyennes non-locales et le modèle TV sont perçus comme des extrêmes dans une décomposition biais-variance de l'erreur. Le chapitre conclut sur quelques considérations sur la loi sur les patches qui atteint l'optimum dans cette décomposition.

Abstract In this chapter we reconsider the Total Variation (TV) prior as a model for images. First we analyze the link between this model and empirical local statistics of images. We show that the restriction of the TV model on patches as little as 2×2 is a good approximation for the empirical distribution of those patches. However the choice of the numerical scheme for the gradient norm turns out to be significant, and the schemes compatible with geometrical properties of total variation (coarea formula, or gradient associated to a divergence operator) do not achieve the maximal adequacy to the patch distribution, which makes statistical and geometrical points of view conflicting. Then in a framework of patch distribution estimation, the Non-Local means and the TV model are viewed as extreme situations in a bias-variance decomposition of the error. The chapter ends with considerations about the patch distribution which achieves the optimal trade-off of this decomposition.

6.1 Introduction

The point of view over Total Variation (TV) minimization has been Bayesian until now. The first chapter can be considered as a discussion on the risk, while Chapter 5 discusses the locality of the model underlying TV minimization. Here in this chapter we discuss the prior model relative to TV-minimization.

The chapter contains two main sections. First Section 6.2 argues in favor of the flexibility of this TV prior. In particular, the geometrical properties of Total Variation, such as its interpretation in terms of cumulated length of level lines (coarea formula), do not hold any more when we change it slightly, whereas its statistical properties can be improved. This gives justification to the following Section 6.3 where the Non-Local means (NL-means) and the TV-denoising are viewed as extreme situations in a bias-variance decomposition of the L^2 -error, the NL-means being associated with a high complexity prior and the TV-denoising with a low complexity prior. We briefly discuss the existence and the complexity of a statistical prior model which achieves the bias-variance trade-off.

6.2 Prior in Total Variation denoising

6.2.1 Gradient norm is a good measure of small patch distribution

Total-variation based denoising methods such as TV minimization [Rudin et al. 1992] or TV-LSE suggest that the images are drawn from the distribution

$$P(u) = \frac{1}{Z} \exp(-\beta TV(u)).$$

This prior model has several good denoising qualities which have been enumerated previously in the thesis. Now one can legitimately wonder whether this prior is relevant with the natural images.

In a local framework the TV prior is a model for the distribution of patches. Let us begin with the smallest possible patches, i.e. 2×2 patches, and derive from natural images a prior on these patches. Let us assume that we have a large database \mathcal{B} of such patches taken from natural images. An estimate of the 2×2 patches distribution is given by a 4-dimensional histogram of \mathcal{B} , which writes for $p = (p_1, p_2, p_3, p_4)$

$$\hat{P}(p) = \frac{1}{|\mathcal{B}|} \sum_{k_1, k_2, k_3, k_4} \sum_{q \in \mathcal{B}} \mathbf{1}_{k_i Q_i \leq q_i < (k_i+1)Q_i, 1 \leq k_i \leq 4} \cdot \mathbf{1}_{k_i Q_i \leq p_i < (k_i+1)Q_i, 1 \leq k_i \leq 4}$$

where (Q_1, Q_2, Q_3, Q_4) denote the dimensions of the hyper-bins. State-of-the-art density estimation allows much richer estimates than histograms, thanks to kernels or wavelets for instance. However the histogram techniques are mature and user-friendly; this is why we confine ourselves to these. The tuning of the bins size (Q_i) is crucial in the estimation of the patches density (see [Birgé and Rozenholc 2006] and references therein). A bias-variance decomposition shows that it should be selected neither too small (the variance would be very high) neither too large (the bias would be large). Considering the 4 components of the patches as equivalent (the 4 gray levels should approximately have the same marginal distribution in the interval $[0, 256)$), the Q_i




Figure 6.1: AIC criterion (6.1) plotted against the bins size. First the Akaike criterion is directly computed on the 2×2 patches of Lena (512×512) (solid curve). The maximum is reached for a bin size of 17, which is quite high compared to the 256 gray levels of the image, but compatible with the fact that Lena’s patches are points in a 4-D space, and do not manage to fill it densely. Then the Akaike criterion was computed on the patches of Lena, assuming rotation and axial symmetry invariance (dashed curve). In other words, each patch of the database is duplicated 8 times, each of the duplicate being a rotated or/and symmetrized version of the original patch. Under these conditions, the maximum is reached for a bin size equal to 15. It could be predicted that it was less than 17 because the 4-D space contains 8 times more points, and a good bin size decreases when the number of points increases (but it decreases slowly).

are taken all equal to Q . This parameter Q is optimized via the AIC criterion of Akaike [1974] which is suitable in dimension more than 1. The AIC criterion consists in assigning to Q the maximizer of

$$\sum_{k \in \mathbb{Z}^4} N_k \log \left(\frac{N_k}{|\mathcal{B}|} \left(\frac{256}{Q} \right)^4 \right) - \left[\left(\frac{256}{Q} \right)^4 - 1 \right] \quad (6.1)$$

where

$$N_k = \sum_{q \in \mathcal{B}} \mathbf{1}_{k_i Q \leq q_i < (k_i+1)Q, 1 \leq i \leq 4}$$

denotes the number of patches in \mathcal{B} belonging to the bin indexed by $k = (k_1, k_2, k_3, k_4)$. The function which should be maximized is the log-likelihood of the histogram whose bins have size Q^4 , penalized by the number $(256/Q)^4$ of bins in the space $[0, 256]^4$.

Figure 6.1 shows the AIC criterion with respect to Q . The solid curve is obtained when the database \mathcal{B} is the set of all the 2×2 patches of the Lena image, and the dashed curve corresponds to the same database plus their $k\pi/2$ -rotated patches and symmetrized patches, in order to achieve rotation and symmetry invariance (hence it contains 8 times more patches than the previous database). The optimal Q is 17 for the raw database containing 511×511 patches,




Figure 6.2: Correlation between $T_\alpha(p)$ (gradient norm scheme) and $-\log(\hat{P}(p))$ with respect to α . Left (Lena’s patches): the maximum correlation (corr ≈ 0.92) is reached for $\alpha \approx -0.24$. Right (Barbara’s patches): the maximum correlation (corr ≈ 0.94) is reached for $\alpha \approx -0.43$. The same experiment carried out on other images always make T_α with a certain $-1 < \alpha < 0$ reach the best correlation. This is a hint in favor of some particular gradient norm schemes which should be more adapted to image modeling than others.

and 15 for the invariant case, which is not so much when dealing with dimension 4. The optimal Q is a slowly decreasing function of the number of patches in \mathcal{B} .

A 4-dimensional histogram which is less dependent of the bins boundaries can be obtained by averaging histograms shifted by all the possible $(t_1, t_2, t_3, t_4) \in [0, Q - 1]^4$. Finally a patch distribution estimate is obtained. Let us compare it with the Total Variation prior on the 2×2 patches

$$P(p) = \frac{1}{Z} \exp(-\beta T_\alpha(p)), \quad (6.2)$$

where $T_\alpha(p)$, defined by

$$T_\alpha(a, b, c, d) = \sqrt{\frac{(a-b)^2 + (b-c)^2 + (c-d)^2 + (d-a)^2 + \alpha[(a-c)^2 + (b-d)^2]}{2 + 2\alpha}},$$

is a consistent scheme for the gradient norm $|\nabla p|$ depending on a parameter α .

$$p = \begin{array}{|c|c|} \hline a & d \\ \hline b & c \\ \hline \end{array}$$

(a, b, c, d) denotes p ’s components in the counterclockwise order beginning, say, with the top-left corner. T_α is defined for $\alpha > -1$, otherwise the expression inside the square-root can become negative.

Figure 6.2 shows the correlation between $T_\alpha(p)$ and $-\log \hat{P}(p)$ with respect to α , and Figure 6.3 the corresponding cloud of points. A correlation equal to 1 would mean that $T_\alpha(p)$ and $-\log \hat{P}(p)$ are perfectly proportional, and that the Total Variation distribution (6.2) on patches would be exact on the database. In the experiments we held, the correlation between $T_\alpha(p)$




Figure 6.3: Clouds of points representing $-\log \hat{P}(p)$ with respect to the gradient norm. Top: the scheme for the gradient norm is T_α , with the α achieving maximal correlation ($\alpha = -0.25$, $\text{corr} \approx 0.92$ for Lena, left, and $\alpha = -0.43$, $\text{corr} \approx 0.94$ for Barbara, right). Second row: the scheme for the gradient norm is $\sqrt{(a-b)^2 + (a-d)^2}$; it is related to a divergence operator, but the direct dependence is less clear than with T_α ($\text{corr} \approx 0.85$ for Lena, $\text{corr} \approx 0.88$ for Barbara). Bottom row: the scheme for the gradient norm is $|a-b| + |a-d|$; it is compatible with the coarea formula, but the direct dependence between the two quantities is not clearer either ($\text{corr} \approx 0.85$ for Lena, $\text{corr} \approx 0.88$ for Barbara again). (These 2-D histograms were passed through a logarithm to have visible contrast).

and $-\log \hat{P}(p)$ finds a maximum for a certain α , and the corresponding correlation is more than 0.9 (corr ≈ 0.92 for Lena, corr ≈ 0.95 for Barbara). This means that the Total Variation distribution (with α not too far from the optimum in T_α) is a good approximation of the patches' distribution.

The simpler gradient norm schemes

$$T(p) = \sqrt{(a-b)^2 + (a-d)^2}, \quad (6.3)$$

adopted in Chambolle [2004] for instance, and the anisotropic scheme

$$T(p) = |a-b| + |a-d|, \quad (6.4)$$

considered in [Darbon and Sigelle 2005; Chambolle 2005], yield smaller correlations (see Figure 6.3).

Numerical experiments have been carried out on several images and it comes that the scheme T_α always reaches a better correlation than T given on (6.3) and (6.4). Furthermore the parameter α which maximizes this correlation has always been found inside the interval $(-1, 0)$. This gives a track for designing a scheme for images gradient norms adapted to image local statistics.

For instance the Total Variation minimization of [Rudin et al. 1992] or TV-LSE (Part 1) carried out with T_α (with $\alpha \in (0, 1)$) is probably slightly better than other schemes, because the prior distribution is more faithful to the real image statistics. The drawback of these schemes is that they do not seem to be associated to any dual divergence operator, contrary to T in (6.3) and (6.4) which allows fast algorithms for Total Variation minimization [Chambolle 2004], unless $\alpha = \infty$. Besides they do not satisfy exactly the coarea formula [Ambrosio et al. 2000], hence the fast graph-cut based algorithm for TV minimization [Darbon and Sigelle 2005] cannot apply. Nevertheless our algorithm for LSE denoising, presented in Chapter 2, works for any scheme for TV satisfying the assumptions stated at the beginning of Chapter 2, and could be run with a scheme consistent with image statistics.

6.2.2 Optimizing the total variation prior

Deriving a global distribution on images from a distribution on its patches is a difficult task even if long-range interactions are neglected, because the patches are not independent. However some information about local features in the images, such as typically the gradient norm, can be assembled thanks to the FRAME algorithm [Zhu et al. 1998] (FRAME standing for *Filters, Random fields And Maximum Entropy*) and a global distribution can be derived. The idea is the following. The global distribution is selected as the one which maximizes the entropy among the distributions matching the available information. The entropy is maximized to make sure no spurious information is brought inside the model.

More precisely, let $(\phi_n(u))_{1 \leq n \leq N}$ be any real functions on images. These functions can be either linear or non-linear. They can typically be local and translation invariant, and can be designed to catch the local features of images. For example, they can be Gabor filters, wavelet coefficients, or any feature testing. Then the FRAME algorithm is based on the maximization of the entropy

$$H(p) = - \sum_{u \in \mathbb{R}^\Omega} p(u) \log p(u)$$

under constraints of the type

$$\forall n \in \{1, \dots, N\}, \quad \mathbb{E}_p[\phi_n(u)] = \sum_{u \in \mathbb{R}^\Omega} \phi_n(u) p(u) = \mu_n,$$

where μ_n is the supposedly known expectation of $\phi_n(u)$, not forgetting that $\sum_{u \in \mathbb{R}^\Omega} p(u) = 1$. The Euler-Lagrange equation related to the optimization problem writes

$$\exists \nu \in \mathbb{R}, \exists \lambda_n \in \mathbb{R}, \forall u \in \mathbb{R}^\Omega, \quad \log p(u) + 1 + \sum_{n=1}^N \lambda_n \phi_n(u) + \nu = 0,$$

which yields

$$\forall u \in \mathbb{R}^\Omega, \quad p(u) = \frac{1}{Z} \exp\left(-\sum_{n=1}^N \lambda_n \phi_n(u)\right)$$

Zhu et al. [1998] prove that if $(\phi_1(u), \dots, \phi_N(u))$ has a positive definite covariance matrix (which holds when they are independent), then there exists at most one solution, which can be computed by solving

$$\frac{d\lambda_n}{dt} = \mathbb{E}_p[\phi_n(u)] - \mu_n \quad \forall n \in \{1, \dots, N\}.$$

In our case more information is available than simple expectations $\mathbb{E}_p[\phi_n(u)]$ alone. In the previous subsection we saw that the gradient norm of a patch is a good indicator of its probability. Assume that the distribution of $|\nabla u|$ is known. We denote it $p_{|\nabla u|}$. We seek a global distribution p which is translation invariant and compatible with the marginal distribution $p_{|\nabla u|}$. As in FRAME, a good choice for p is the one which maximizes the entropy, subject to the constraints

$$\forall x \in \Omega, \forall y \in \mathbb{R}, \quad \sum_u \mathbf{1}_{|\nabla u(x)|=y} p(u) = p_{|\nabla u|}(y) \quad \text{and} \quad \sum_u p(u) = 1.$$

The Lagrangian formulation for this maximization problem amounts to the minimization of

$$\sum_{u \in \mathbb{R}^\Omega} p(u) \log p(u) + \sum_x \sum_y \lambda_x(y) \sum_u \mathbf{1}_{|\nabla u(x)|=y} p(u) + \nu \sum_u p(u).$$

The minimizer then satisfies

$$\forall u, \quad p(u) = \frac{1}{Z} \exp\left(-\sum_{x \in \Omega} \lambda_x(|\nabla u(x)|)\right).$$

If p is translation invariant, the values (λ_x) do not depend on x , and then there exists a mapping λ , called potential (in relation to Gibbs fields), such that

$$\forall u, \quad p(u) = \frac{1}{Z} \exp\left(-\sum_{x \in \Omega} \lambda(|\nabla u(x)|)\right).$$

It can be iteratively computed by the gradient descent

$$\frac{d\lambda(y)}{dt} = \frac{1}{\Omega} \sum_{x \in \Omega} \mathbb{E}_p[\mathbf{1}_{|\nabla u(x)|=y}] - p_{|\nabla u|}(y).$$

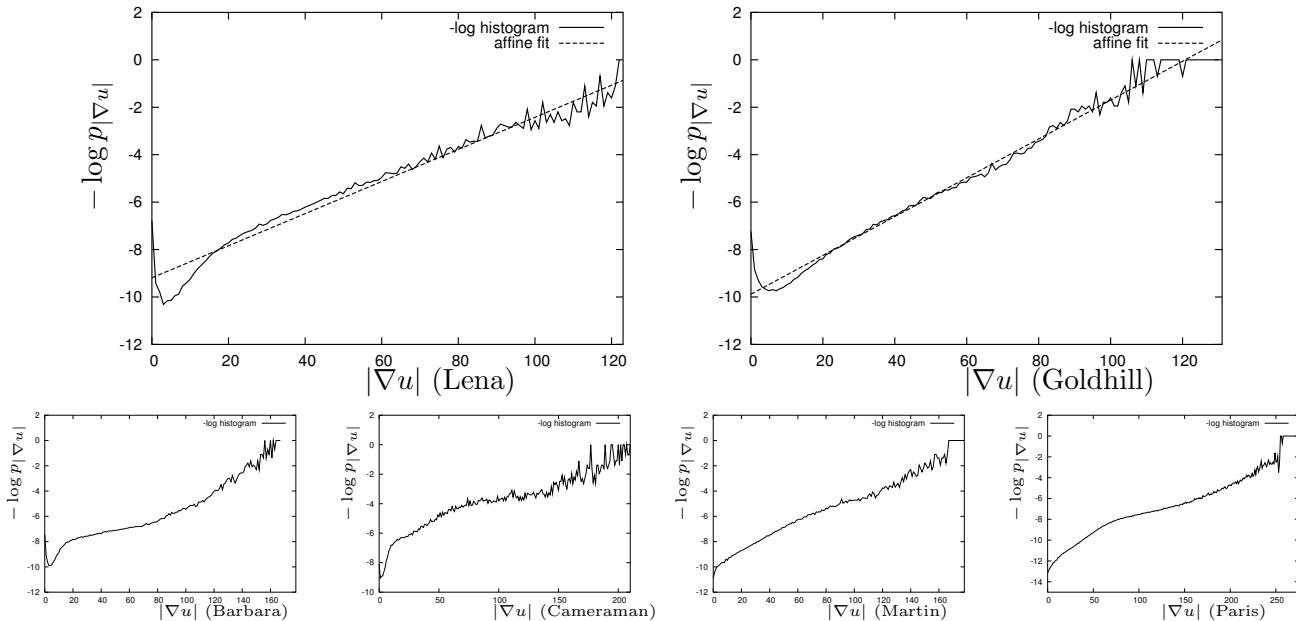


Figure 6.4: Distribution of gradient norm for different images. Here the opposite of logarithm of the histograms are plotted. Apart from a particular behavior for small gradient norms, the curves can be practically fitted with affine lines. It means that the histograms are not far from approximating an exponential distribution. The scheme for gradient norm was chosen to be $T_{-1/2}$, which is a good generic value for high dependence between gradient norm and $-\log P(\text{patch})$. The last two images from which the histograms are computed (Martin and Paris) are shown on Figure 6.5.

Notice that when λ is linear the distribution p is the total variation distribution. However this case do not seem to be strictly compatible with the gradient norm distribution estimated on natural images.

Experiments The gradient norm distribution $p_{|\nabla u|}$ on natural images is not far from being exponential. Indeed, Figure 6.4 plots the opposite logarithm of gradient norm histograms of several images and shows that the curves are approximately affine. However it can be refined into a distribution written

$$p_{|\nabla u|}(y) = \frac{1}{Z} \exp(-\mu(y))$$

where μ is a smooth function which is approximately linear except for small values of the gradient norm, for which μ generally twists along an inflection point and attains a strict minimum. The behavior of μ for high values of gradient norms is difficult to examine because the samples contains few of these.

Our adaptation of FRAME algorithm was first run in the case where $p_{|\nabla u|}$ is approximated by an exponential distribution

$$p_{|\nabla u|}(y) = l e^{-ly} \mathbf{1}_{y>0}.$$



Figure 6.5: Data images used in Figure 6.4. Martin is 800×533 , it comes from the scan of a traditional silver-based photograph, and contains important natural noise. Paris is 2592×1944 and comes from an amateur digital camera.

The algorithm is quite heavy to run because at each iteration of the gradient descent we need to simulate images drawn from distribution p . It comes that the potential λ in the global distribution, shown on Figure 6.6 (left) is a smooth and concave function which mainly contains two affine parts. The smaller gradient norms correspond to a large slope, while larger gradient norm correspond to a small slope in λ .




Figure 6.6: (a) Potential λ obtained by FRAME algorithm for exponential target distribution. (b) Potential λ obtained by FRAME algorithm for Lena's gradient norm statistics target (the histogram was regularized by a Gaussian kernel with s.d. 1 to get convergence more easily).

Afterwards the adapted FRAME algorithm was run on an empirical histogram of gradient norms, computed on image Lena. The histogram is very irregular and the convergence of FRAME algorithm was difficult to obtain, so the choice was made to regularize a little bit the histogram, thanks to convolution by a Gaussian kernel with s.d. 1. The potential λ is shown on Figure 6.6 (right) and has a minimum in the small gradient norms, much more contrasted as in

the histograms shown on Figure 6.4, and becomes almost flat for larger gradient norms.

The potentials λ obtained in our experiments are not exactly linear, even if a linear function could be a very rough approximation of the exponential distribution case (left). This implies that the Total Variation may be improved with respect to the gradient statistics of images.

It seems then that the prior

$$p(u) = \frac{1}{Z} \exp\left(-\sum_x \lambda(|\nabla u(x)|)\right)$$

with λ looking as one of the curves of Figure 6.6 should then be more realistic than the Total Variation prior. Several authors [Blake and Zisserman 1987; Bouman and Sauer 1993; Nikolova 2005] already noticed that the minimization of

$$\|u - v\|^2 + 2\sigma^2 \sum_x \lambda(|\nabla u(x)|) \quad \text{with} \quad \lambda(y) = \min\{\alpha y^2, \beta\} \quad (\alpha, \beta > 0)$$

was a denoiser with some very interesting properties, such as the possible recovery of edges and smooth regions together. This corresponds to the MAP denoising using the potential λ in the prior.

The obtained potential can also be linked to the “regularized Total Variation” techniques, where authors such as [Blomgren et al. 1997] propose to minimize energies of the form

$$\|u - v\|^2 + 2\sigma^2 \sum_x \lambda(|\nabla u(x)|) \quad \text{with} \quad \lambda(y) = y^{g(y)}$$

where g is decreasing and satisfies

$$\begin{cases} g(y) \xrightarrow{y \rightarrow 0} 2 \\ g(y) \xrightarrow{y \rightarrow +\infty} 1. \end{cases}$$

The corresponding potential $\lambda(y)$ is non-convex, and can be understood as a more realistic potential for natural images.

Of course all these potentials are non-convex. Several methods for the minimization of possibly non-convex energies are available, such as Metropolis sampler, simulated annealing [Geman and Geman 1987], graduated non-convexity [Blake and Zisserman 1987] or more recently stochastic continuation [Robini et al. 2007], but yet the MAP estimates are more difficult to obtain in the non-convex case. Nevertheless this study gives arguments for the fact that Total Variation denoising mainly works because of its statistical interpretation. The geometric properties of Total Variation, such as its interpretation in terms of cumulated length of the image’s level lines (coarea formula), are not enough to explain the denoising performances; the associated statistical model is a good approximation of image prior.

This gives motivation to make the prior model move away from the TV model.

6.3 Bias-variance trade-off in the prior complexity

6.3.1 Bayesian framework for NL-means

As the TV minimization, the Non-Local means (NL-means) algorithm, proposed recently by Buades, Coll and Morel in their seminal paper [Buades et al. 2005] can be seen as a statistical denoising method. The risk is L^2 and gives a Least Square Error estimate. Besides the prior model is built from the patches of the noisy image.

Indeed, the denoising operator NL , applied on an image $v \in \mathbb{R}^\Omega$, is defined by

$$\forall x \in \Omega, NL(v)(x) = \frac{1}{Z_x} \sum_{y \in \Omega} e^{-\frac{\|v(\mathcal{N}_x) - v(\mathcal{N}_y)\|_2^2}{2\sigma^2}} v(y), \quad \text{where } Z_x = \sum_{y \in \Omega} e^{-\frac{\|v(\mathcal{N}_x) - v(\mathcal{N}_y)\|_2^2}{2\sigma^2}}$$

as described in [Buades et al. 2005] (the only subtlety we omit here is the choice of a weighted Euclidean norm over the patches, which is the equivalent to the smooth windows for local TV-MAP in Chapter 5 Section 5.3.2). Z_x is the normalizing factor, σ^2 is the noise's variance and \mathcal{N} is a neighborhood shape with $\mathcal{N}_x = x + \mathcal{N}$. The gray level assigned to pixel x is a weighted average of all the gray levels $v(y)$ of the image, and the weights measure a certain similarity between the neighborhood of x and the neighborhood of the other pixels y .

Let us define a prior on the patches $p \in \mathbb{R}^{\mathcal{N}}$ from those of the image v . A simple histogram is given by

$$P(p) = \frac{1}{|\Omega|} \sum_{y \in \Omega} \mathbf{1}_{v(\mathcal{N}_y)}(p).$$

Assume that the image v is a noisy version of an image u , that is $v = u + n$ with n being a white Gaussian noise. The Bayesian estimate $\hat{u}(x)$ of $u(x)$ with L^2 -risk and prior P on the patches is given by the posterior expectation

$$\hat{u}(x) = \mathbb{E}_P[u(x)|v(\mathcal{N}_x)]$$

which develops into

$$\hat{u}(x) = \frac{1}{Z} \sum_{p \in \mathbb{R}^{\mathcal{N}}} P(p) e^{-\frac{\|v(\mathcal{N}_x) - p\|_2^2}{2\sigma^2}} p(0)$$

where $p(0)$ denotes the central gray value of the patch p . Now using our definition for the prior P , and switching the sums, we get

$$\hat{u}(x) = \frac{1}{Z'} \sum_{y \in \Omega} \sum_{p \in \mathbb{R}^{\mathcal{N}}} \mathbf{1}_{v(\mathcal{N}_y)}(p) e^{-\frac{\|v(\mathcal{N}_x) - p\|_2^2}{2\sigma^2}} v(y).$$

For a given y , a $p \in \mathbb{R}^{\mathcal{N}}$ yielding a non-zero term can only be $p = v(\mathcal{N}_y)$, and finally,

$$\hat{u}(x) = NL(v)(x).$$

NL-means is therefore a LSE Bayesian estimation, and the prior model for NL-means is hence based on the histogram of the noisy patches taken on the image.

6.3.2 Bias and variance for TV-denoising and for NL-means

The prior distribution in Total variation denoising

$$P_{TV}(u) = \frac{1}{Z} e^{-\beta TV(u)}$$

is global, but a localization procedure as in Chapter 5 shows that P_{TV} is also a prior on patches. This prior is based on a very poor statistics. Indeed only one quantity computed on an image or a patch (the total variation) determines its probability; longer range interactions are omitted, and two images having possibly very different geometrical properties can have the same total variation. Conversely the prior distribution used in the Non-Local means denoising

$$P_{NL}(p) = \frac{1}{|\Omega|} \sum_{y \in \Omega} \mathbf{1}_{v(\mathcal{N}_y)}(p)$$

is a local model on patches, and is based on the exhaustive statistics of the patches, and the patches contain noise. The prior distributions P_{TV} and P_{NL} , related to these two denoising methods, have therefore a complexity lying at the antipodes each other, the first one with an extremely low complexity, and the other with large complexity. They can be interpreted as extreme situations in a bias-variance decomposition. Indeed, let P be a prior and \hat{u}_P a Bayes estimate of u based on the prior P . Then the expected L^2 -error with respect to the noise n can be decomposed into a bias and a variance, following

$$\begin{aligned} \mathbb{E}_n \|\hat{u}_P - u\|^2 &= \mathbb{E}_n \|(\mathbb{E}_n \hat{u}_P - u) + (\hat{u}_P - \mathbb{E}_n \hat{u}_P)\|^2 \\ &= \|\mathbb{E}_n \hat{u}_P - u\|^2 + \mathbb{E}_n \|\hat{u}_P - \mathbb{E}_n \hat{u}_P\|^2, \end{aligned}$$

where the first term stands for the squared bias of the estimate \hat{u}_P , and the second term for its variance. High complexity priors usually achieve low bias rates but the variance can be high because the estimate is very sensitive to the noise. Instead, low complexity priors usually imply large bias because the model cannot fit the data enough while the variance is small. Here for the NL-means, it should be all the more true that the prior is based on the statistics of *noisy* patches: the variance of the estimator should be quite high.

Even if the TV-MAP denoising and the NL-means are not associated to the same risk, their bias and variance are compared in Figure 6.7. Bias and variance were approximated thanks to averages over 50 noise samples. They were computed on two details of Lena and Barbara. The detail from Lena (around the eyes) is typically an image that is not very well treated by NL-means because the patches of the eyes have very few similar patches around them, and the denoising is much less efficient there than on the cheek for instance. The estimate is then much more dependent of the noise. This is an archetype for a locally large variance in NL-means. Conversely, the detail from Barbara (the thin stripes) is a typical image where TV denoising is not efficient, because the oscillating patterns have high frequency and are viewed by Total Variation as noise. This is essentially due to the non-adaptability in the TV model, leading to a large bias.

Is it then possible to find a prior which achieves a bias-variance trade-off? This is briefly and partially discussed in the following subsection.

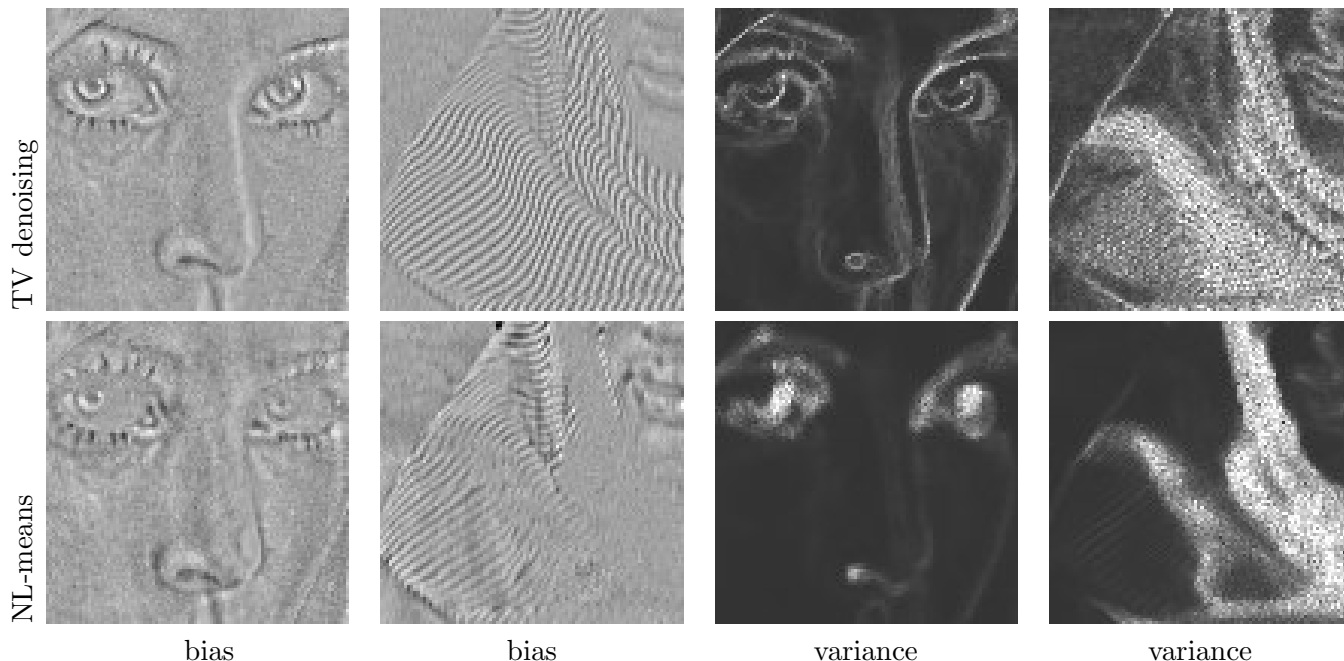


Figure 6.7: Bias-variance decomposition of the estimation error in TV-MAP denoising and in NL-means. Two pieces of images are selected. On the 4 left images, the bias ($u - \mathbb{E}_n \hat{u}_P$) is plotted in the case of TV denoising and NL-means (the expectation was approximated thanks to an average over 50 noises). In the first image (Lena's eyes), the bias is concentrated on the edges. For TV denoising virtually all the structure of the image can be found in the bias. For NL-means the bias is more uniform. In the second image (striped fabric in Barbara), again, the bias contains lesser structure in NL-means than in TV-denoising. On the 4 right images, the variance $\mathbb{E}_n |\hat{u}_P - \mathbb{E}_n \hat{u}_P|^2$ is plotted (again the expectation was approximated thanks to an average over 50 noises). In the case of NL-means denoising, the variance is concentrated on the eyes of Lena, and on the turbulent parts of the striped fabric, whereas in TV denoising, the variance is more uniform. In that sense we can say that TV denoising is more biased, and that NL-means leads to higher variances.

6.3.3 Argument for the existence of a non-trivial optimum

Several small hints indicate that there should exist a non-trivial optimum in the bias-variance decomposition, i.e. that should be different from the TV prior and from the NL-means prior.

A first argument is pointed out in [Azzabou et al. 2007], which can be interpreted as an attempt for complexity reduction in the NL-means prior. Namely the similarity weights in the NL-means

$$\omega_{x,y} = \frac{1}{Z_x} e^{-\frac{\|v(\mathcal{N}_x) - v(\mathcal{N}_y)\|^2}{2\sigma^2}}$$

were found too much dependent on the noise. The authors propose a principal component analysis on the patches of the noisy image, in order to select the components having the largest associated eigenvalues. It happens that these components are much more robust to the noise than the ones associated to low eigenvalues. New weights are then proposed, based on the components

of the patches with large eigenvalues only. This leads to an estimate which is expected to be more robust to the noise. Some denoising examples are shown where this technique is better than NL-means.

A second argument is the following. Assume that the patches of natural images can be decomposed into independent components (for instance after an independent component analysis), such that their global distribution writes as a product of marginal distributions

$$\forall p = (p_1, \dots, p_N) \in \mathcal{X}^N, \quad f(p) = \prod_{k=1}^N f_k(p_k)$$

with f_k probability distribution over a finite set \mathcal{X} . Here we aim at estimating f as precisely as possible from simple histograms of patches sampled from f . We allow the histograms \hat{f} to inject only some components $k \in \{1, \dots, N\}$, that is, for a sample of patches $\mathcal{P} = \{p^1, \dots, p^{|\mathcal{P}|}\}$,

$$\forall x \in \mathcal{X}^N, \quad \hat{f}_{A, \mathcal{P}}(x) = \frac{1}{|\mathcal{X}|^{N-|A|}} \prod_{k \in A} \frac{1}{|\mathcal{P}|} \sum_{i=1}^{|\mathcal{P}|} \mathbf{1}_{p_k^i}(x_k),$$

where $A \subset \{1, \dots, N\}$. A represents the indices of components which are taken into account in the histogram. It amounts to consider that the f_k with $k \notin A$ are approximated by a uniform distribution. Its cardinality $|A|$ reflects the model complexity in the estimate. Intuitively, in average, a histogram $\hat{f}_{A, \mathcal{P}}$ will correctly estimate f if A captures the most *informative* components of the patches. Conversely, if A contains too many non-informative components, and if $|\mathcal{P}|$ is quite small, then the histogram $\hat{f}_{A, \mathcal{P}}$ will deviate from f , because a uniform distribution is harder to approximate with histograms on finite samples than informative distributions, such as very peaked distributions. Besides, if the sample size $|\mathcal{P}|$ is assumed to be very large, but if the sample is noisy, then similarly, the components containing little information will be overwhelmed with noise and should be thrown away (as in [Azzabou et al. 2007]) and the more informative components should be kept, to yield an efficient estimation for f . In both cases (small sample size, or noisy samples), it seems that an optimal A can be reached, with $1 < |A| < N$, containing the indices of the most informative f_k . A can be viewed as the set which achieves a bias-variance trade-off in the estimation error.

Of course this program is feasible if we are able to predict or guess through the sample which components are informative.

What we meant by *informative* distribution was not defined on purpose. It is highly dependent on the distance we consider between f and $\hat{f}_{A, \mathcal{P}}$. For example, if we aim at minimizing $\|\hat{f}_{A, \mathcal{P}} - f\|^2$ in average, it is provable and logical that the information criterion over f_k is its L^2 squared norm, which satisfies

$$\|f_k\|_2^2 = \sum_{x \in \mathcal{X}} f_k(x)^2 = \mathbb{P}_{x, y \sim f_k}(x = y).$$

$\|f_k\|_2^2$ measures some information in a distribution, in the sense that uniform distribution has lowest $\|f_k\|_2^2$ and Dirac the highest. We do not doubt that another criterion of proximity between f and $\hat{f}_{A, \mathcal{P}}$ should lead to other information measures. For instance the Kullback-Leibler distance $D(\hat{f}_{A, \mathcal{P}} \| f)$ should probably lead to a classical entropy-based information measure.

We have no general result proving the existence of a non-trivial complexity prior reaching an optimal error, either in the model prior estimation or directly in the denoising task. However we believe that recent works on patch statistics such as [Lee et al. 2003; Huang and Mumford 1999b; Huang et al. 2000] could help finding “informative” components on patches of natural images, able to increase the complexity of the TV prior and to reach less biased models. Simultaneously, applying invariance of patch statistics under simple geometric operations, such as symmetry or rotation, could decrease the complexity of NL-means prior.

Conclusion

In this thesis we focused on several image denoising models based on the total variation (TV), derived from transpositions to some aspects of the Non-Local means (NL-means) algorithm. This was made possible thanks to a common Bayesian framework for Total Variation minimization and for the Non-Local means.

The first aspect we were interested in was the choice of the Bayesian risk. The adaptation of TV minimization to a least square error risk as in the NL-means has led us to a denoising method which was analyzed in the first part of the thesis. Motivations for this change of risk were given in the first chapter: the new risk was expected to imply fewer distortions in the output local statistics than with the “hit-or-miss” risk attached to TV minimization. A stochastic algorithm based on MCMCs was found especially well suited to the implementation of the method. Convergence was proven, but a precise convergence criterion was needed because the high dimensionality of the images implied quite slow convergence. Subject to the introduction of other parameters, all being justified and automatically tuned thanks to simple algorithms derived from recent Statistics papers, we elaborated an efficient stopping criterion. Thanks to it, the unpleasant impacts of a slow convergence were contained, and the denoising method could then be computable. Several interesting properties of the denoising method were proven. In particular it was shown to induce no staircasing effect; this is the main contribution of this method compared to the TV minimization, and it confirms that the least square error risk leads to fewer distortions than the “hit-or-miss” risk. Numerous tests were carried out; the results were very satisfying for several reasons: first the staircasing is indeed absent from the denoised images; secondly the contrast of some isolated pixels, very often enhanced by TV minimization, is reduced with our method; last but not least, the visual aspect of the denoised images is much more natural with our method than with TV minimization.

The second area of interest which was examined in this thesis deals with the locality in the denoising. Namely the NL-means had been noticed to be more efficient when “localizing” it, i.e. when disregarding far away pixels in the denoising of a pixel. This localization procedure was carried out in the TV minimization framework, leading to a TV-based neighborhood filter. An algorithm was provided as well as the partial differential equation associated to a regularized version of this neighborhood filter, which happens to be simply the heat equation. The local denoising is necessarily associated to a certain size of neighborhood; the optimal size may be understood as a bias-variance trade-off. Two versions of this local TV-based denoising were explored; the first one was associated to a Lagrangian formulation, and the other to a constrained formulation, the latter giving rise to a somewhat automatically adaptive denoising. No serious attempt was made to “localize” the TV denoising operator defined with a least square error risk

because of its numerical complexity, but it seems that it should be a very nice hybrid denoiser between TV-minimization and NL-means.

The last missing link between TV denoising and NL-means lies in the definition of an image prior model. The last chapter questions the total variation model as a suitable model for discrete images. In particular the selection of the numerical scheme for the gradient norm happens to be crucial for the model's adequacy to local statistics of natural images. Besides, a global model is derived from the local statistics of the gradient norm, and experiments show that this model is a slight deviation from the Total Variation model. This allows us to consider modified versions of the TV model. Actually the TV model has a very low "complexity" compared to the model underlying the NL-means which has high complexity, and a bias-variance decomposition argues in favor of some optimal prior model with a complexity lying in between those of TV model and NL-means.

However a prior model that fits reality as perfectly as possible is not necessarily a good prior in practice. Indeed in a general Bayesian estimation framework, a prior model is always combined with some risk that has to be minimized. The minimization of this risk leads to an estimate that may induce some distortions. These distortions really bias the estimation accuracy, and even perfect priors may induce imperfect estimation. They occur with Maximum A Posteriori (MAP) estimation but also for other risks, such as the L^2 -risk.

We would like to point out here that conversely, prior models could be distorted with respect to the risk in order to reduce the bias of estimates. For example a feature which gets swelled in the estimation can be voluntarily squeezed in the model, in order to obtain unbiased estimates. Several facts hint at a distortion of the prior model in favor of the rare configurations, i.e. the model used in estimation should probably give higher probability to events having low probability in the reality-sticking model, at the expense of frequent events whose probability should be reduced.

A first argument for this is the example of [Nikolova 2007] where a signal simulated from the 1-dimensional TV model is corrupted with an impulse noise or Laplace noise. A Laplace model for the noise is used in the MAP denoising procedure in both cases. Results are much more convincing on the effective impulse noise than on the Laplace noise. The Laplace distribution has a wider "spread" than the impulse noise distribution, and the distribution with its Dirac at zero is shrunk down to a Lebesgue measurable density, while the probability of other events gets swelled.

Another argument comes from the results of Chapter 3, Section 3.5, where a Least-Square-Error (LSE) estimate is shown under some conditions to be equivalent to a MAP estimate, provided that the prior distributions p_{LSE} (resp. p_{MAP}) in the LSE denoiser (resp. MAP denoiser) satisfy

$$p_{MAP} \approx p_{LSE} * G_{\sigma},$$

where G_{σ} is a Gaussian kernel. Under the assumption that the LSE estimate induces lesser distortions than the MAP estimate, we can deduce that a prior model should rather be smoother than the realistic one in the context of a MAP estimation. Again rare events are enhanced and frequent ones are reduced.

This does not hold for MAP estimation only. Indeed a very simple case shows that the same

phenomenon can also occur in LSE estimation. Let us describe it here. Let us assume that samples are drawn from a real distribution

$$P_\tau(x) = \tau\delta_{x_1}(x) + (1 - \tau)\delta_{x_2}(x), \quad x \in \mathbb{R}^n$$

where n is assumed to be a “large” integer, $0 < \tau < 1/2$, and $x_1 \neq x_2$. P_τ is very similar to a Bernoulli distribution, except from the space dimension. Let us assume too that we only have one sample of P_τ which is corrupted by a white Gaussian noise, namely

$$y = x + n, \quad \text{where } x \sim P_\tau \text{ and } n \sim G_\sigma.$$

Let us assume that we estimate x via a LSE risk based on another “Bernoulli” prior model P_t , the denoising prior, with $0 < t < 1$. Then

$$\widehat{x_1 + n} = \frac{e^{-\frac{\|n\|^2}{2\sigma^2}}tx_1 + e^{-\frac{\|x_1+n-x_2\|^2}{2\sigma^2}}(1-t)x_2}{e^{-\frac{\|n\|^2}{2\sigma^2}}t + e^{-\frac{\|x_1+n-x_2\|^2}{2\sigma^2}}(1-t)} \quad \text{and} \quad \widehat{x_2 + n} = \frac{e^{-\frac{\|x_2+n-x_1\|^2}{2\sigma^2}}tx_1 + e^{-\frac{\|n\|^2}{2\sigma^2}}(1-t)x_2}{e^{-\frac{\|x_2+n-x_1\|^2}{2\sigma^2}}t + e^{-\frac{\|n\|^2}{2\sigma^2}}(1-t)}.$$

As the space has high dimensionality we make the following approximation

$$\|x_2 + n - x_1\|^2 \approx \|x_1 - x_2\|^2 + \|n\|^2$$

which is a valid approximation for most $n \sim G_\sigma$. Then denoting

$$\alpha = e^{-\frac{\|x_1-x_2\|^2}{2\sigma^2}},$$

the estimates $\widehat{x_1 + n}$ and $\widehat{x_2 + n}$ get simpler, and the average error writes

$$\mathbb{E}_{n,x} \|\widehat{x + n} - x\|_2^2 \approx \left[\frac{\tau(1-t)^2}{(t + \alpha(1-t))^2} + \frac{(1-\tau)t^2}{(\alpha t + (1-t))^2} \right] \cdot \alpha^2 \|p_2 - p_1\|^2 =: \mathcal{E}(t)$$

This function \mathcal{E} of $t \in (0, 1)$ can be optimized. Actually \mathcal{E} can be proven to be convex, and its derivative \mathcal{E}' to satisfy

$$\mathcal{E}'(\tau) < 0 \quad \text{and} \quad \mathcal{E}'(1/2) > 0.$$

Then there exists a denoising prior P_t that minimizes the average error, and such that the corresponding t satisfies

$$\tau < t < \frac{1}{2}$$

(recall that P_τ is the real prior model from which the samples are drawn, and P_t is the artificial prior used in the denoising procedure). The best prior associated to the LSE risk, here selected among “Bernoulli” priors, again enhances the rarer event and decreases the importance of the other one.

No result was found yet proving that this phenomenon can be generalized. In particular an example with log-concave prior and L^2 -risk would be welcome. Of course in realistic cases the rare event should probably not be all equally enhanced. The information brought by the denoising prior would then be lost, and the denoising quality would decrease. In image denoising,

certain configurations should probably be enhanced, such as visually important features, like T-junctions, corners, and so on, whereas other rare events should probably be reduced. In any case it seems that the prior should be designed in connection to the risk used in the estimation, and that it is generally not equal to the realistic model.

These considerations on the model prior can also be transposed to the noise prior. Indeed, in most denoising methods, in TV minimization e.g., the method noise contains structures stolen from the image, even if the noise is modeled as white and Gaussian noise. It means that the method noise seriously deviates from the noise model, which initially allows no significant structure.

The duality between image estimation and noise estimation, when dealing with symmetric risks (Proposition 1.3), implies that the noise prior should also be designed relatively to the risk and the expected distortions. In this case, avoiding the distortions is probably easier to model. This is the idea of *a contrario* methods [Desolneux et al. 2008] which assume that the noise is easier to model than real-life features. Indeed if the noise is really white and Gaussian, then the probability of a n -uple of pixels should not depend on the relative locations of these pixels. This property could be inserted in the denoising prior in the form of constraints or not. It could lead to method noises containing lesser structure. In the same time, the denoised image is likely to be of better quality because information does not go away from the image in favor of the method noise.

Another point is the link between the locality and the selection of the risk. [Rue 1995; Rue and Hurn 1997] propose risk functions associated to more or less localized estimates. Besides, image quality measures closer to human vision criterions than SNR or PSNR (see [Wang et al. 2004] for instance) can probably be associated to a risk. These non-conventional risks used with Total Variation or the prior implied in the NL-means may yield lesser visual distortions.

These remarks hold for image denoising, but may also apply to other image-based estimation tasks, such as interpolation, dequantization and other restoration issues.

Bibliography

- Akaike, H. (1974). A new look at the statistical model identification. *IEEE Trans. Automatic Control*, 19:716–723.
- Almansa, A., Ballester, C., Caselles, V., and Haro, G. (2008). A TV based restoration model with local constraints. *J. Sci. Comput.*, 34(3):209–236.
- Alter, F., Caselles, V., and Chambolle, A. (2005a). Evolution of characteristic functions of convex sets in the plane by the minimizing total variation flow. *Interfaces Free Bound.*, 7(1):29–53.
- Alter, F., Durand, S., and Froment, J. (2005b). Adapted total variation for artifact free decomposition of JPEG images. *J. Math. Imaging Vis.*, 23(2):199–211.
- Alvarez, L., Guichard, F., Lions, P.-L., and Morel, J.-M. (1993). Axioms and fundamental equations of image processing. *Archives for Rational Mechanics*, 123:199–257.
- Ambrosio, L., Fusco, N., and Pallara, D. (2000). *Functions of bounded variation and free discontinuity problems*. Oxford Mathematical Monographs. The Clarendon Press Oxford University Press, New York.
- Andreu, F., Ballester, C., Caselles, V., and Mazòn, J. M. (2000). Minimizing total variation flow. *Comptes Rendus de l’Académie des Sciences*, 331(11):867–872.
- Andrieu, C. and Moulines, É. (2006). On the ergodicity properties of some adaptive MCMC algorithms. *Ann. Appl. Probab.*, 16(3):1462–1505.
- Arulampalam, S., Maskell, S., and Gordon, N. (2002). A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Transactions on Signal Processing*, 50:174–188.
- Aubert, G. and Kornprobst, P. (2006). *Mathematical Problems in Image Processing: Partial Differential Equations and the Calculus of Variations (second edition)*, volume 147 of *Applied Mathematical Sciences*. Springer-Verlag.
- Aujol, J.-F., Aubert, G., Blanc-Féraud, L., and Chambolle, A. (2005). Image decomposition into a bounded variation component and an oscillating component. *J. Math. Imaging Vis.*, 22(1):71–88.

- Aujol, J.-F. and Chambolle, A. (2005). Dual norms and image decomposition models. *Int. J. Comput. Vision*, 63(1):85–104.
- Aujol, J.-F., Gilboa, G., Chan, T., and Osher, S. (2006). Structure-texture image decomposition—modeling, algorithms, and parameter selection. *Int. J. Comput. Vision*, 67(1):111–136.
- Awate, S. and Whitaker, R. (2005). Higher-order image statistics for unsupervised, information-theoretic, adaptive, image filtering. In *CVPR05*, pages II: 44–51.
- Awate, S. and Whitaker, R. (2006). Unsupervised, information-theoretic, adaptive image filtering for image restoration. *PAMI*, 28(3):364–376.
- Azzabou, N., Paragios, N., and Guichard, F. (2007). Image denoising based on adapted dictionary computation. *IEEE International Conference on Image Processing*, 3:109–112.
- Azzabou, N., Paragios, N., and Guichard, F. (2008). Uniform and textured regions separation in natural images towards MPM adaptive denoising. *Scale Space and Variational Methods in Computer Vision*, 4485:418–429.
- Bédard, M. (2007). Weak convergence of Metropolis algorithms for non-i.i.d. target distributions. *Ann. Appl. Probab.*, 17:1222–1244.
- Bertalmio, M., Caselles, V., Rougé, B., and Solé, A. (2003). TV based image restoration with local constraints. *J. Sci. Comput.*, 19(1-3):95–122.
- Besag, J. (1974). Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society, Series B*, 36(2):192–236.
- Besag, J. (1986). On the statistical analysis of dirty pictures. *Journal of the Royal Statistical Society*, 48(3):259–302.
- Besag, J. (1989). Digital image processing : Towards Bayesian image analysis. *Journal of Applied Statistics*, 16(3):395–407.
- Besag, J., Green, P., Higdon, D., and Mengersen, K. (1995). Bayesian computation and stochastic systems. *Statistical Science*, 10(1):3–41.
- Birgé, L. and Rozenholc, Y. (2006). How many bins should be put in a regular histogram. *ESAIM Probab. Stat.*, 10:24–45 (electronic).
- Blake, A. and Zisserman, A. (1987). *Visual Reconstruction*. MIT Press.
- Blanchet, G. (2006). *Etude des artefacts de flou, ringing et aliasing en imagerie numérique. Application à la restauration*. PhD thesis, ENS Cachan.
- Blomgren, P., Chan, T., Mulet, P., and Wong, C. (1997). Total variation image restoration: numerical methods and extensions. *ICIP*, 3:384–387.
- Borghi, A., Darbon, J., Peyronnet, S., Chan, T. F., and Osher, S. (2008). A simple compressive sensing algorithm for parallel many-core architectures. CAM Report 08-64, UCLA.

- Bouman, C. and Sauer, K. (1993). A generalized Gaussian image model for edge-preserving MAP estimation. *IEEE Trans. on Image Processing*, 2:296–310.
- Breyer, L. A. and Roberts, G. O. (2000). From Metropolis to diffusions: Gibbs states and optimal scaling. *Stochastic Process. Appl.*, 90(2):181–206.
- Brezis, H. R. (1965-1966). Les opérateurs monotones. *Séminaire Choquet, Initiation à l'analyse*, 5(2).
- Brox, T. and Weickert, J. (2006). A TV flow based local scale estimate and its application to texture discrimination. *Journal of Visual Communication and Image Representation*, 17(5):1053–1073.
- Buades, A., Coll, B., Lisani, J.-L., and Sbert, C. (2007). Conditional image diffusion. *Inverse Problems and Imaging*, 1(4):593–608.
- Buades, A., Coll, B., and Morel, J.-M. (2005). A review of image denoising algorithms, with a new one. *Multiscale Model. Simul.*, 4(2):490–530 (electronic).
- Buades, A., Coll, B., and Morel, J.-M. (2006a). Neighborhood filters and PDE's. *Numer. Math.*, 105(1):1–34.
- Buades, A., Coll, B., and Morel, J.-M. (2006b). The staircasing effect in neighborhood filters and its solution. *IEEE Transactions on Image Processing*, 15(6):1499–1505.
- Buades, A., Coll, B., and Morel, J.-M. (2008). Nonlocal image and movie denoising. *Int. J. Comput. Vision*, 76(2):123–139.
- Buccigrossi, R. W. and Simoncelli, E. P. (1999). Image compression via joint statistical characterization in the wavelet domain. *IEEE Transactions on Image Processing*, 8(12):1688–1701.
- Caffarelli, L. A. and Cabré, X. (1995). *Fully Nonlinear Elliptic Equations*. AMS Bookstore.
- Caselles, V., Chambolle, A., and Novaga, M. (2007). The discontinuity set of solutions of the TV denoising problem and some extensions. *Multiscale Model. Simul.*, 6(3):879–894.
- Catté, F., Dibos, F., and Koepfler, G. (1995). A morphological scheme for mean curvature motion and applications to anisotropic diffusion and motion of level sets. *SIAM J. Numer. Anal.*, 32(6):1895–1909.
- Chalmond, B. (2003). *Modeling and Inverse Problems in Image Analysis*. Applied Mathematical Science 155. Springer Verlag, New York, Inc.
- Chambolle, A. (2004). An algorithm for total variation minimization and applications. *J. Math. Imaging Vision*, 20(1-2):89–97. Special issue on mathematics and image analysis.
- Chambolle, A. (2005). Total variation minimization and a class of binary MRF models. In *EMMCVPR*, pages 136–152.

- Chambolle, A. and Lions, P.-L. (1997). Image recovery via total variation minimization and related problems. *Numer. Math.*, 76(2):167–188.
- Chan, T., Marquina, A., and Mulet, P. (2000). High-order total variation-based image restoration. *SIAM J. Sci. Comput.*, 22(2):503–516.
- Chan, T. and Wong, C. (1998). Total variation blind deconvolution. *IEEE Transactions on Image Processing*, 7(3):370–375.
- Chan, T. F., Esedoglu, S., and Park, F. E. (2007). Image decomposition combining staircase reduction and texture extraction. *J. Vis. Comun. Image Represent.*, 18(6):464–486.
- Coifman, R. R. and Donoho, D. L. (1995). Translation-invariant de-noising. *Lecture notes in statistics: Wavelets and statistics*, pages 125–150.
- Combettes, P. L. and Pesquet, J. C. (2004). Image restoration subject to a total variation constraint. *IEEE Transactions on Image Processing*, 13(9):1213–1222.
- Combettes, P. L. and Wajs, V. R. (2004). Theoretical analysis of some regularized image denoising methods. In *ICIP04*, pages II: 969–972.
- Combettes, P. L. and Wajs, V. R. (2005). Signal recovery by proximal forward-backward splitting. *Multiscale Modeling and Simulation*, 4(4):1168–1200.
- Criminisi, A., Perez, P., and Toyama, K. (2004). Region filling and object removal by exemplar-based image inpainting. *IEEE Transactions on Image Processing*, 13(9):1200–1212.
- Csiszar, I. and Talata, Z. (2006). Consistent estimation of the basic neighborhood of Markov random fields. *The Annals of Statistics*, 34(1):123–145.
- Dabov, K., Foi, A., Katkovnik, V., and Egiazarian, K. (2006). Image denoising with block-matching and 3D filtering. In *Electronic Imaging'06, Proc. SPIE 6064*, volume 30.
- Darbon, J. and Sigelle, M. (2005). A fast and exact algorithm for total variation minimization. In *IbPRIA (1)*, pages 351–359.
- Darbon, J. and Sigelle, M. (2006a). Image restoration with discrete constrained total variation Part I: Fast and exact optimization. *J. Math. Imaging Vis.*, 26(3):261–276.
- Darbon, J. and Sigelle, M. (2006b). Image restoration with discrete constrained total variation Part II: Levelable functions, convex priors and non-convex cases. *J. Math. Imaging Vis.*, 26(3):277–291.
- Desolneux, A., Moisan, L., and Morel, J.-M. (2008). *From Gestalt Theory to Image Analysis: a probabilistic approach*, volume 34 of *Interdisciplinary Applied Mathematics*. Springer Verlag.
- Dobson, D. C. and Santosa, F. (1996). Recovery of blocky images from noisy and blurred data. *SIAM J. Appl. Math.*, 56(4):1181–1198.

- Durand, S. and Froment, J. (2003). Reconstruction of wavelet coefficients using total variation minimization. *SIAM J. Sci. Comput.*, 24(5):1754–1767 (electronic).
- Durand, S. and Nikolova, M. (2007). Denoising of frame coefficients using l^1 data-fidelity term and edge-preserving regularization. *Multiscale Model. Simul.*, 6(2):547–576 (electronic).
- Durrett, R. (1996). *Probability: theory and examples*. Duxbury Press, Belmont, CA, second edition.
- Efros, A. A. and Leung, T. K. (1999). Texture synthesis by non-parametric sampling. In *ICCV '99: Proceedings of the International Conference on Computer Vision-Volume 2*, page 1033, Washington, DC, USA. IEEE Computer Society.
- Evans, L. C. and Gariepy, R. F. (1992). *Measure theory and fine properties of functions*. Studies in Advanced Mathematics. CRC Press, Boca Raton, FL.
- Federer, H. (1969). *Geometric measure theory*. Springer-Verlag.
- Fessler, J. A. (1996). Mean and variance of implicitly defined biased estimators (such as penalized maximum likelihood): Applications to tomography. *IEEE Tr. Im. Proc.*, 5:493–506.
- Fort, G., Moulines, E., Roberts, G. O., and Rosenthal, J. S. (2003). On the geometric ergodicity of hybrid samplers. *J. Appl. Probab.*, 40(1):123–146.
- Fu, H., Ng, M. K., Nikolova, M., and Barlow, J. L. (2006). Efficient minimization methods of mixed l_2 - l_1 and l_1 - l_1 norms for image restoration. *SIAM Journal on Scientific Computing*, 27(6):1881–1902.
- Fu, H., Ng, M. K., Nikolova, M., Barlow, J. L., and Ching, W.-K. (2005). Fast algorithms for l_1 norm/mixed l_1 and l_2 norms for image restoration. In Gervasi, O., Gavrilova, M. L., Kumar, V., Laganà, A., Lee, H. P., Mun, Y., Taniar, D., and Tan, C. J. K., editors, *ICCSA (4)*, volume 3483 of *Lecture Notes in Computer Science*, pages 843–851. Springer.
- Gelman, A., Roberts, G. O., and Gilks, W. R. (1996). Efficient Metropolis jumping rules. In *Bayesian statistics, 5 (Alicante, 1994)*, Oxford Sci. Publ., pages 599–607. Oxford Univ. Press, New York.
- Geman, D. and Koloydenko, A. (1999). Invariant statistics and coding of natural microimages. In *CVPR 99*.
- Geman, D. and Yang, C. (1995). Nonlinear image recovery with half-quadratic regularization. *IEEE Transactions on Image Processing*, 5(7):932–946.
- Geman, S. and Geman, D. (1987). Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. In *Readings in computer vision: issues, problems, principles, and paradigms*, pages 564–584. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- Geyer, C. (1992). Practical Markov Chain Monte Carlo. *Statistical Science*, 7(4):473–483.

- Gilboa, G., Sochen, N., and Zeevi, Y. Y. (2003). Texture preserving variational denoising using an adaptive fidelity term. pages 137–144.
- Gilboa, G., Sochen, N., and Zeevi, Y. Y. (2006). Variational denoising of partly-textured images by spatially varying constraints. *IEEE Transactions on Image Processing*, 15(8):2281–2289.
- Goldfarb, D. and Yin, W. (2005). Second-order cone programming methods for total variation-based image restoration. *SIAM J. Sci. Comput.*, 27(2):622–645.
- Gordon, N., Salmond, D., and Smith, A. F. M. (1993). Novel approach to non-linear and non-Gaussian Bayesian state estimation. *Radar and Signal Processing, IEEE Proceedings-F*, 140(2):107–113.
- Gousseau, Y. and Morel, J.-M. (2001). Are natural images of bounded variation? *SIAM J. Math. Anal.*, 33(3):634–648 (electronic).
- Grimmett, G. R. and Stirzaker, D. R. (2001). *Probability and random processes*. Oxford University Press, New York, third edition.
- Guichard, F. and Malgouyres, F. (1998). Total variation based interpolation. In *Proceedings of Eusipco'98*, volume 3, pages 1741–1744.
- Guichard, F., Moisan, L., and Morel, J.-M. (2002). A review of PDE models in image image processing and image analysis. *Journal de Physique IV*, 12:137–154.
- Guichard, F. and Morel, J.-M. (2001). Image analysis and PDE's.
- Hammersley, J. M. and Clifford, P. (1971). Markov field on finite graphs and lattices. (Unpublished).
- Hastings, W. K. (1970). Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109.
- Hida, T. and Nomoto, H. (1964). Gaussian measure on the projective limit space of spheres. *Proc. Japan Acad.*, 40(5):301–304.
- Hiriart-Urruty, J. and Lemaréchal, C. (1993). *Convex analysis and minimization algorithms: A review of PDE models in image image processing and image analysis*, volume 305-306. Springer Verlag.
- Hochbaum, D. S. (2001). An efficient algorithm for image segmentation, Markov random fields and related problems. *J. ACM*, 48(4):686–701.
- Huang, G., Lee, A. B., and Mumford, D. (2000). Statistics of range images. In *Proceedings of CVPR*, pages 541–547, South Carolina.
- Huang, G. and Mumford, D. (1999a). Image statistics for the British aerospace segmented database. MTPC preprint, Brown University.

- Huang, G. and Mumford, D. (1999b). The statistics of natural images and models. In *Proceedings of IEEE Comp. Vision and Pattern Recognition*, pages 541–547.
- Jalobeanu, A., Blanc-Féraud, L., and Zerubia, J. (1998). Estimation d’hyperparamètres pour la restauration d’images satellitaires par une méthode MCMCML. Research Report 3469, INRIA.
- Jarner, S. F. and Hansen, E. (2000). Geometric ergodicity of Metropolis algorithms. *Stochastic Process. Appl.*, 85(2):341–361.
- Kervrann, C. and Boulanger, J. (2006). Optimal spatial adaptation for patch-based image denoising. *IEEE Trans. on Image Processing*, 15(10):2866–2878.
- Kervrann, C. and Boulanger, J. (2007). Local adaptivity to variable smoothness for exemplar-based image denoising and representation. *International Journal of Computer Vision*.
- Kervrann, C., Boulanger, J., and Coupé, P. (2007). Bayesian non-local means filter, image redundancy and adaptive dictionaries for noise removal. In *Proc. Conf. Scale-Space and Variational Meth. (SSVM’ 07)*, pages 520–532, Ischia, Italy.
- Kindermann, S., Osher, S., and Jones, P. W. (2005). Deblurring and denoising of images by nonlocal functionals. *Multiscale Model. Simul.*, 4(4):1091–1115 (electronic).
- Kipnis, C. and Varadhan, S. R. S. (1986). Central limit theorem for additive functionals of reversible Markov processes and applications to simple exclusions. *Comm. Math. Phys.*, 104(1):1–19.
- Koloydenko, A. A. (2000). *Modeling natural microimage statistics*. PhD thesis, University of Massachusetts. Director-Donald Geman.
- Koloydenko, A. A. (2004). Adaptive coding of microstructure of natural images. *Eurandom seminar*.
- Lafon, S., Keller, Y., and Coifman, R. R. (2006). Data fusion and multicue data matching by diffusion maps. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(11):1784–1797.
- Lee, A. B., Pedersen, K. S., and Mumford, D. (2003). The nonlinear statistics of high-contrast patches in natural images. *Int. J. Comput. Vision*, 54(1-3):83–103.
- Levine, S., Chen, Y., and Stanich, J. (2004). Image restoration via nonstandard diffusion. Technical Report 04-01, Dept mathematics and computer science, Duquesne University.
- Luo, B., Aujol, J.-F., Gousseau, Y., Ladjal, S., and Maître, H. (2007). Resolution-independent characteristic scale dedicated to satellite images. *IEEE Transactions on Image Processing*, 16(10):2503–2514.
- Malgouyres, F. (2002). Minimizing the total variation under a general convex constraint for image restoration. *IEEE Transactions on Image Processing*, 11(12):1450–1456.

- Malgouyres, F. (2007). Rank related properties for basis pursuit and total variation regularization. *Signal Process.*, 87(11):2695–2707.
- Mallat, S. (1998). *A wavelet tour of signal processing*. Academic Press Inc., San Diego, CA.
- Maurey, B. (2003). Inégalité de Brunn-Minkowski-Lusternik, et autres inégalités géométriques et fonctionnelles. *Séminaire Bourbaki*, 2003/2004(Astérisque No. 299, Exp. No. 928):vii, 95–113.
- Mengersen, K. L. and Tweedie, R. L. (1996). Rates of convergence of the Hastings and Metropolis algorithms. *Ann. Statist.*, 24(1):101–121.
- Meyer, Y. (2001). *Oscillating patterns in image processing and nonlinear evolution equations*, volume 22 of *University Lecture Series*. American Mathematical Society, Providence, RI. The fifteenth Dean Jacqueline B. Lewis memorial lectures.
- Meyn, S. and Tweedie, R. (1993). *Markov Chains and Stochastic Stability*. Springer Verlag, London.
- Moisan, L. (1998). Affine plane curve evolution: a fully consistent scheme. *IEEE Transactions on Image Processing*, 7(3):411–420.
- Moisan, L. (2001). Extrapolation de spectre et variation totale pondérée. *GRETSI'01*.
- Moisan, L. (2007). How to discretize the total variation of an image? *ICIAM'07*.
- Moreau, J. (1962). Fonctions convexes duales et points proximaux dans un espace hilbertien. *C.R. Acad. Sci. Paris Sér. A Math.*, 255:2897–2899.
- Moreau, J. (1965). Proximité et dualité dans un espace hilbertien. *Bulletin de la Société Mathématique de France*, 93:273–299.
- Mumford, D. and Gidas, B. (2001). Stochastic models for generic images. *Quart. Appl. Math.*, 59(1):85–111.
- Neal, P. and Roberts, G. (2006). Optimal scaling for partially updating MCMC algorithms. *Ann. Appl. Probab.*, 16(2):475–515.
- Nikolova, M. (1997). Estimées localement fortement homogènes. *Compte-rendus de l'académie des sciences*, t. 325:665–670.
- Nikolova, M. (2000). Local strong homogeneity of a regularized estimator. *SIAM J. Appl. Math.*, 61(2):633–658 (electronic).
- Nikolova, M. (2004). Weakly constrained minimization: application to the estimation of images and signals involving constant regions. *J. Math. Imaging Vision*, 21(2):155–175.
- Nikolova, M. (2005). Analysis of the recovery of edges in images and signals by minimizing nonconvex regularized least-squares. *Multiscale Model. Simul.*, 4(3):960–991 (electronic).

- Nikolova, M. (2007). Model distortions in Bayesian MAP reconstruction. *Inverse Probl. Imaging*, 1(2):399–422.
- Olshausen, B. and Simoncelli, E. (2001). Natural image statistics and neural representation. *Annu. Rev. Neurosci.*, 24:1193–1216.
- Ordentlich, E., Seroussi, G., Verdú, S., Weinberger, M. J., and Weissman, T. (2003). A discrete universal denoiser and its application to binary images. In *ICIP (1)*, pages 117–120.
- Osher, S., Solé, A., and Vese, L. (2003). Image decomposition and restoration using total variation minimization and the H^{-1} norm. *Multiscale Model. Simul.*, 1(3):349–370 (electronic).
- Ouvrard, J.-Y. (2000). *Probabilités 2*. Cassini.
- Owens, J. D., Luebke, D., Govindaraju, N., Harris, M., Kruger, J., Lefohn, A. E., and Purcell, T. J. (2007). A survey of general-purpose computation on graphics hardware. *Computer Graphics Forum*, 26(1):80–113.
- Perona, P. and Malik, J. (1990). Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Anal. Mach. Intell.*, 12(7):629–639.
- Peyré, G. (2008). Image processing with non-local spectral bases. *to appear in SIAM Multiscale Modeling and Simulation*.
- Popat, K. and Picard, R. (1997). Cluster-based probability model and its application to image and texture processing. *IEEE Transactions on Image Processing*, 6(2):268–284.
- Portilla, J. and Simoncelli, E. P. (2000). A parametric texture model based on joint statistics of complex wavelet coefficients. *Int. J. Comput. Vision*, 40(1):49–70.
- Prékopa, A. (1973). On logarithmic concave measures and functions. *Acta Scientiarum Mathematicarum*, 34:335–343.
- Ring, W. (2000). Structural properties of solutions of total variation regularization problems. *ESSAIM, Math Modelling and Numerical Analysis*, 34:799–840.
- Robert, C. (1996). *Méthodes de Monte Carlo par chaînes de Markov*. Statistique Mathématique et Probabilité. [Mathematical Statistics and Probability]. Éditions Économica, Paris.
- Roberts, G. O. and Rosenthal, J. S. (1997). Geometric ergodicity and hybrid Markov chains. *Electron. Comm. Probab.*, 2(2):13–25 (electronic).
- Roberts, G. O. and Rosenthal, J. S. (1998). Two convergence properties of hybrid samplers. *Ann. Appl. Probab.*, 8(2):397–407.
- Roberts, G. O. and Rosenthal, J. S. (2001). Optimal scaling for various Metropolis-Hastings algorithms. *Statist. Sci.*, 16(4):351–367.
- Roberts, G. O. and Tweedie, R. L. (1996). Geometric convergence and central limit theorems for multidimensional Hastings and Metropolis algorithms. *Biometrika*, 83(1):95–110.

- Robini, M. C., Lachal, A., and Magnin, I. E. (2007). A stochastic continuation approach to piecewise constant reconstruction. *IEEE Transactions on Image Processing*, 16(10):2576–2589.
- Roussas, G. G. (1990). Nonparametric regression estimation under mixing conditions. *Stochastic Process. Appl.*, 36(1):107–116.
- Rousseeuw, P. and Leroy, A. (1987). *Robust regression and outlier detection*. John Wiley & Sons, Inc., New York, NY, USA.
- Ruderman, D. (1994). The statistics of natural images. *Network*, 5(4):517–548.
- Rudin, L. I., Osher, S., and Fatemi, E. (1992). Nonlinear total variation based noise removal algorithms. *Phys. D*, 60(1-4):259–268.
- Rue, H. (1995). New loss functions in Bayesian imaging. *Journal of the American Statistical Association*, 90:900–908.
- Rue, H. and Hurn, M. (1997). Loss functions for Bayesian image analysis. *Scandinavian Journal of Statistics*, 24:103–114.
- Savage, J. and Chen, K. (2006). On multigrids for solving a class of improved total variation based PDE models. In *Proceedings of the 1st International Conference PDE-based Image Processing and Related Inverse Problems*, CMA Oslo. Springer Verlag.
- Simoncelli, E. (1999). Bayesian denoising of visual images in the wavelet domain. *Lecture Notes in Statistics*, 41.
- Smith, S. M. and Brady, J. M. (1997). SUSAN—a new approach to low level image processing. *Int. J. Comput. Vision*, 23(1):45–78.
- Srivastava, A., Lee, A. B., Simoncelli, E. P., and Zhu, S.-C. (2003). On advances in statistical modeling of natural images. *J. Math. Imaging Vision*, 18(1):17–33. Special issue on imaging science (Boston, MA, 2002).
- Steidl, G., Weickert, J., Brox, T., Mrázek, P., and Welk, M. (2004). On the equivalence of soft wavelet shrinkage, total variation diffusion, total variation regularization, and SIDes. *SIAM J. Numer. Anal.*, 42(2):686–713 (electronic).
- Strong, D. and Chan, T. (1996). Exact solutions to total variation regularization problems. CAM Report 96-41, UCLA Math Department.
- Strong, D. and Chan, T. (2003). Edge-preserving and scale-dependent properties of total variation regularization. *Inverse Problems*, 19(6):S165–S187. Special section on imaging.
- Strong, D. M., Blomgren, P., and Chan, T. F. (1997). Spatially adaptive local feature-driven total variation minimizing image restoration. In *Proceedings of SPIE*, pages 222–233.
- Tierney, L. (1994). Markov chains for exploring posterior distributions. *Ann. Statist.*, 22(4):1701–1762. With discussion and a rejoinder by the author.

- Tropp, J. A. (2006). Just relax: convex programming methods for identifying sparse signals in noise. *IEEE Trans. on Inform. Theory*, 52(3):1030–1051.
- Vese, L. A. and Osher, S. J. (2003). Modeling textures with total variation minimization and oscillating patterns in image processing. *Journal of Scientific Computing*, 19:553–572.
- Wang, Z., Bovik, A. C., Sheikh, H. R., and Simoncelli, E. P. (2004). Image quality assessment: from error visibility to structural similarity. *IEEE Trans. on Image Processing*, 13(4):600–612.
- Yaroslavsky, L. and Eden, M. (1996). *Fundamentals of Digital Optics*. Birkhauser, Boston.
- Zhu, S. C. and Mumford, D. (1997). Prior learning and Gibbs reaction-diffusion. *IEEE Trans. on Pattern Anal. Mach. Intell.*, 19(11):1236–1250.
- Zhu, S. C., Wu, Y., and Mumford, D. (1998). Filters, random fields and maximum entropy (FRAME): Towards a unified theory for texture modeling. *Int. J. Comput. Vision*, 27(2):107–126.

Résumé: Le modèle ROF (Rudin, Osher, Fatemi), introduit en 1992 en utilisant la variation totale comme terme de régularisation pour la restauration d’images, a fait l’objet de nombreuses recherches théoriques et numériques depuis. Dans cette thèse, nous présentons de nouveaux modèles inspirés de la variation totale mais construits par analogie avec une méthode de débruitage beaucoup plus récente et radicalement différente : les moyennes non locales (*NL-means*). Dans une première partie, nous transposons le modèle ROF dans un cadre bayésien, et montrons que l’estimateur associé à un risque quadratique (moyenne a posteriori) peut être calculé numériquement à l’aide d’un algorithme de type MCMC (*Monte Carlo Markov Chain*), dont la convergence est soigneusement contrôlée compte tenu de la dimension élevée de l’espace des images. Nous montrons que le débruiteur associé permet notamment d’éviter le phénomène de *staircasing*, défaut bien connu du modèle ROF. Dans la deuxième partie, nous proposons tout d’abord une version localisée du modèle ROF et en analysons certains aspects : compromis biais-variance, EDP limite, pondération du voisinage, etc. Enfin, nous discutons le choix de la variation totale en tant que modèle a priori, en confrontant le point de vue géométrique (modèle ROF) au cadre statistique (modélisation bayésienne).

Mots-clés: Débruitage d’images, variation totale, modèles bayésiens, maximum a posteriori, moyenne a posteriori, effet de *staircasing*, algorithme de *Monte-Carlo Markov Chains*, filtre à voisinage, moyennes non-locales, compromis biais-variance.

Abstract: The ROF (Rudin, Osher, Fatemi, 1992) model, introducing the total variation as regularizing term for image restoration, has since been dealt with intense numerical and theoretical research. In this thesis we present new models inspired by the total variation but built by analogy with a much more recent method and diametrically opposed to it: the non-local means. In the first part we transpose the ROF model into a Bayesian framework, and show that the estimator associated to a quadratic risk (posterior expectation) can be numerically computed thanks to a MCMC (Monte Carlo Markov Chain) algorithm, whose convergence is carefully controlled, considering the high dimensionality of the image space. We notably prove that the associated denoiser avoids the staircasing effect, a well-known artefact that frequently occurs in ROF denoising. In the second part of the thesis we first propose a localized version of the ROF model, and analyze several aspects: bias-variance trade-off, limiting PDE, neighborhood weighting, etc. Last but not least we reconsider the choice of total variation as prior image model, by setting the geometrical point of view (ROF model) against the statistical framework (Bayesian modeling).

Keywords: Image denoising, total variation, Bayesian models, maximum a posteriori, posterior expectation, staircasing effect, Monte-Carlo Markov Chains, neighborhood filter, non-local means, bias-variance trade-off.