



HAL
open science

Stability and Regions of Attraction of Image-Based Visual Servoing Systems

Alessandro Colotti

► **To cite this version:**

Alessandro Colotti. Stability and Regions of Attraction of Image-Based Visual Servoing Systems. Automatic. École centrale de Nantes, 2023. English. NNT : 2023ECDN0046 . tel-04546588

HAL Id: tel-04546588

<https://theses.hal.science/tel-04546588v1>

Submitted on 15 Apr 2024

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.

MEMOIRE DE DOCTORAT DE

L'ÉCOLE CENTRALE DE NANTES

ÉCOLE DOCTORALE N° 602
Sciences de l'Ingénierie et des Systèmes
Spécialité : *Automatique*

Par

Alessandro COLOTTI

Stability and Regions of Attraction of Image-Based Visual Servoing Systems

Projet de recherche doctoral présenté et soutenu à l'École Centrale de Nantes, le 21/12/2023
Unité de recherche : UMR 6004, Laboratoire de Sciences du Numérique de Nantes (LS2N)

Rapporteurs avant soutenance :

Sophie TARBOURIECH Directrice de recherche CNRS, CNRS Toulouse
Philippe MARTINET Directeur de recherche, INRIA Côte d'Azur

Composition du Jury :

Président :	François CHAUMETTE	Directeur de recherche, INRIA Rennes
Examineurs :	Sophie TARBOURIECH	Directrice de recherche CNRS, CNRS Toulouse
	Philippe MARTINET	Directeur de recherche, INRIA Côte d'Azur
	Nicolas ANDREFF	Professeur des universités, Université de Franche-Comté
	Denis EFIMOV	Directeur de recherche, INRIA Lille
	Franck PLESTAN	Professeur des universités, École Centrale de Nantes
Dir. de recherches doctorales :	Sébastien BRIOT	Directeur de recherche CNRS, École Centrale de Nantes
Co-enc. de recherches doctorales :	Alexandre GOLDSZTEJN	Chargé de recherche CNRS, École Centrale de Nantes

ACKNOWLEDGEMENT

This Ph.D. has been a unique and incredibly rewarding experience, for which I am extremely grateful. Starting out, I simply thought it would be a job like many others; what I found, instead, was an experience that changed my life forever, during which I had so much fun that every now and then I even wondered how it was possible that I even got paid to do it. First of all, I want to thank the people who made it possible by offering me the opportunity to come and stay here in Nantes.

Tout d'abord, je tiens à remercier Alexandre, pour m'avoir guidé dans ce monde totalement nouveau pour moi et pour avoir toujours été là quand j'en avais besoin. Avec toi, j'ai appris le plaisir et la joie de faire de la recherche, sans se prendre au sérieux mais en ne perdant jamais de vue la rigueur et le respect que la recherche exige. C'est surtout grâce à toi que je peux me croire chercheur aujourd'hui, et que je suis persuadé qu'en fin de compte, on fait bien de la poésie.

À l'autre extrême, je remercie Olivier de m'avoir montré la valeur d'une approche pragmatique à la recherche. Là où j'avais tendance à me perdre dans des envolées pindariques, c'est toi qui m'as ramené sur terre, en me montrant que le *quick and dirty*, même s'il n'est pas très élégant, fonctionne très souvent. Avec Alexandre, vous avez représenté un véritable "yin" et "yang" de la supervision, m'offrant des aperçus de points de vue très différents qui m'ont donné une vue d'ensemble inestimable.

Bien que n'étant pas directement impliqué dans ma supervision, je tiens à remercier Sébastien, à la fois pour ton soutien et parce que c'est grâce à toi que je suis venu en France en premier lieu. De même, je tiens à remercier François pour avoir été une figure de référence et pour m'avoir donné l'occasion de rester plus longtemps dans ce domaine. Je me réjouis de travailler avec toi.

Je tiens à remercier chaleureusement tous les membres du jury : Nicolas Andreff, Denis Efimov, Franck Plestan, et en particulier les rapporteurs, Philippe Martinet et Sophie Tarbouriech, pour m'avoir fait l'honneur d'évaluer mon travail avec autant d'attention et pour m'avoir fourni des commentaires constructifs qui m'ont permis de l'affiner et de l'enrichir.

Je tiens également à remercier Jorge Garcia Fontan et Mohab Safey El Din du LIP6

à Sorbonne Université. Sans votre expertise en algèbre computationnelle, mon travail n'aurait été qu'un château construit sur les nuages.

But the supervisors and permanent staff are not the only ones who have accompanied me on this journey, and I cannot refrain from thanking all the comrades who have made it special. First of all, I want to thank Minh-Quan, for being such a dear friend that it seems impossible that we have known each other for just three years. I also thank the *squadra azzurra* (Vincenzo, Federico and Andrea), Chaolun, Durgesh and Ivan as well, for all the time we spent together and, I hope, for all the time we will still spend together in the future.

I thank all ARMEN Ph.D. students, and especially the open space fellows - Charlotte, Farhan, Hugo, Mathilde, Peng, Selma, and, above all, the dear coffee machine in our office, an honorary member of the team. I also thank all the OGRE (and not only) Ph.D. students whom I got to know at the *Faculté des Sciences et des Techniques* of Nantes Université, and especially Rémi, the two Thomas and Mathieu.

Infine, voglio ringraziare tutte le persone grazie alle quali sono potuto arrivare in Francia in primo luogo. Ringrazio i miei genitori e mio fratello Andrea per l'amore e il supporto, sia materiale che non, che mi hanno sempre dato senza nessuna esitazione. Voglio inoltre ringraziare il mio fratello spirituale Calza, che conosco da quasi venticinque anni, assieme al quale ho affrontato tutte le fasi più importanti della vita di un ragazzo che diventa adulto, costruendo pezzo per pezzo, e con non pochi errori, la bussola morale che mi ha condotto qui oggi, e col quale ho condiviso tante esperienze da riempirci due vite intere. Ringrazio il carissimo Gregorio, che realizzo ora conoscere da più di dieci anni, per essere uno dei miei punti di riferimento in questo mondo contraddittorio, così brutto ma così magnifico, per le innumerevoli discussioni che abbiamo avuto e per i tuoi punti di vista che così spesso mi hanno fatto cambiare idea. È grazie a voi che oggi posso credere di essere diventato una persona migliore, una persona che mai pensavo che sarei potuto essere.

Nonostante tutto, ringrazio anche Barnaba, a casa del quale ho passato una grossa parte della mia adolescenza. Forse è a causa tua se oggi sono una persona peggiore di quanto potrei essere, ma quello che abbiamo condiviso ha creato un legame così forte e profondo che devo ammettere, onestamente, che tornando indietro rifarei tutto. Ringrazio Paolo e Claudio, due ragazzi bellissimi con cui ho fatto amicizia un po' per caso, ma grazie al quale i miei ultimi anni di studi saranno sempre i più belli di tutta la mia carriera accademica. Assieme siamo riusciti a portare a termine idee che sembravano

irrealizzabili, creando un sodalizio che permane tutt'oggi e, spero, che rimarrà per sempre. Infine ringrazio tutta quella brutta ma bella gente col quale sono cresciuto: Lenny, il Pose, il Top e tutti gli altri, dei quali porterò sempre una traccia nella mia anima.

Tengo per ultimo il ringraziamento più importante. Voglio ringraziare, con tutto il mio cuore, mia moglie Martina, per essere sempre stata la colonna portante del mio (e nostro) avvenire e il fascio di luce sul quale ho potuto e posso fare riferimento nei momenti più bui e difficili. Senza di te, la mia vita colerebbe ovunque come un liquido senza contenitore, sparendo nel terreno per perdere ogni significato. Se oggi posso dirmi fiero di chi sono, è principalmente grazie a te, e, grazie anche alla nostra piccola Alice, sono impaziente di vedere il nostro futuro assieme.

TABLE OF CONTENTS

Introduction	9
Contributions	16
Thesis' structure	20
1 Equilibria and stability of energy-decreasing systems	23
1.1 Gradient systems in image-based visual servoing	24
1.2 Dynamical equivalence of generic energy-decreasing systems	32
1.2.1 Fixed points	33
1.2.2 Stable and unstable manifolds	37
1.3 Importance of unstable equilibria: the connectivity tree	42
1.3.1 State-of-the-art and its limits	43
1.3.2 Sublevel sets' connected components	45
1.3.3 The connectivity tree: definition and examples	48
1.3.4 The connectivity tree: formal justification	52
1.3.5 Detection and discrimination of regions of attraction	54
1.4 Non-compact sublevel sets in the visual servoing context	58
1.4.1 Non-compactness of sublevel sets due to unboundedness	59
1.4.2 Non-compactness of sublevel sets due to the state-space's topology	61
2 Computation of the image-based visual servoing energy's critical points	67
2.1 Mathematical models for the visual servoing equilibrium condition	68
2.1.1 Camera state-space representation	69
2.1.2 Extended features representation	71
2.2 Pose reconstruction	72
2.2.1 Reconstruction algorithm	73
2.2.2 Full solution set of the orthogonal Procrustes' problem	74
2.3 Verified computation of all critical points	78
2.3.1 Generalities on computational geometry	80
2.3.2 Symmetry breaking in the coplanar case	82

TABLE OF CONTENTS

3	Representative test cases in image-based visual servoing	85
3.1	Test cases overview	87
3.2	Test case 1, non-planar configuration	90
3.3	Test case 2, planar quadrilateral	95
3.4	Test case 3, from rectangle to square	101
4	Analysis and energy-decrease enforcement for other known controllers	109
4.1	Desired pose approximation controllers	110
4.2	Representative test cases for desired pose approximation	111
4.2.1	Test case 1, 4 points non-planar configuration	114
4.2.2	Test case 2, 5 points non-planar configuration	117
4.2.3	Test case 3, planar square	120
4.3	Meta-controller to enforce energy decrease of an arbitrary control law . . .	121
4.3.1	Illustrative example	123
4.3.2	General definition of the meta-controller	126
4.3.3	Continuity of the meta-controller	127
4.4	Meta-controller in image-based visual servoing	130
5	Handling uncertainties	141
5.1	Practical stability of perturbed dynamical systems	143
5.1.1	Value function definition	144
5.1.2	Continuity of the value function	146
5.1.3	Application to linear tracking systems with bounded velocity target	150
5.2	Extended convergence condition for the super-twisting algorithm	157
5.2.1	Extension to a time-varying integral gain	160
5.2.2	Extension to uncertain input gain	162
5.2.3	Comparison to the state-of-the-art conditions	163
	Conclusion and perspectives	165
	A Proof of Theorem 1.1	169
	Bibliography	173
	Author's publications	173
	In review	173
	Bibliography	174

INTRODUCTION

Visual servoing is a mature area whose formalism is well established [23]. One of its main aims is the control of a robot using the visual information coming from a camera, which is typically mounted at the end effector of the robot. Classically, we assume that the camera has six degrees of freedom and it is controllable in velocity, meaning that it is free to move in any direction and with any angular velocity at any given moment. Additional constraints, like joint limits and the robot’s workspace, as well as inertial effects, are typically handled in a second phase, which is out of scope of this thesis.

The objective of all visual servoing controllers is to minimize an error \mathbf{e} , which is typically defined as:

$$\mathbf{e} = \mathbf{s}(\mathbf{x}) - \mathbf{s}^*, \quad (0.1)$$

where $\mathbf{x} \in \text{SE}(3)$ is the state of the camera and $\mathbf{s}(\mathbf{x})$ are the so-called *visual features*, representing the information that we can gather from the images coming in from the camera. There is a rich literature on the choice of the visual features [19, 23, 28]. Classically, we distinguish between *pose-based* visual servoing, where we first reconstruct the pose of the camera and use directly this information in the control loop, and *image-based* visual servoing, where we use directly the image information (for instance, the Cartesian coordinates of some markers in the screen) to design the controller, but other, hybrid approaches can be found as well.

Even if visual servoing has led to many successful applications and benefits in practice from a large convergence domain and a large robustness to modeling and calibration errors, its stability analysis is still an open theoretical issue when all the six degrees of freedom of the system have to be controlled.

In this regard, pose-based visual servoing [100, 103] seems to be appealing since using a minimal representation of pose as input of the control scheme allows the system to be globally asymptotically stable (i.e., the system will converge whatever its initial configuration) in case the pose is assumed to be perfectly estimated all along the servo. Without speaking about the visibility constraint that imposes a sufficient number of image measurements are always available to estimate the pose, the assumption of perfect pose estimations is very strong and can be violated, as exhibited in [20] for the case of a

simple planar object. This is due to the fact that going from 2D image measurements to pose is an inverse problem which can be ill-conditioned.

The situation is the same for the hybrid and 2 1/2 D approaches [21, 63, 68] that use once again a minimal number of visual features as inputs of the control scheme, some expressed in the image, others in 3D. In that case, global asymptotic stability can be demonstrated by assuming the 3D data involved are perfectly estimated.

Another approach consists in using as visual features the image coordinates of at least three points and their depth [16, 72]. It has been demonstrated in [92] that four equilibrium poses exist in that case, one corresponding to the desired pose and the three others corresponding to unstable equilibria, which means that the system will move away from them as soon as noise will be introduced. Very recently, a deeper study of this scheme has been performed in [34], leading to the nice result that the system is almost globally asymptotically stable in perfect conditions, but also that there exists a set of particular poses that will converge towards these unstable equilibria, meaning they correspond to saddle points.

Very few results exist concerning image-based visual servoing. This approach has the strong advantage of being the more robust to measurement noise since the inputs of the closed-loop control scheme only rely on image measurements. Even if depth is involved in the control scheme since it appears in the interaction matrix from which the control scheme is designed, noise on depth will affect the transient phase of the dynamic system (that is, the trajectory to reach the goal will be perturbed by noise or coarse approximation on depth), but it will have no effect on the accuracy reached at the goal (if the system converges to that goal!). However, up-to-now, only local asymptotic stability has been demonstrated for image-based visual servoing [23], which means one is sure that the system converges if the initial configuration is in the neighborhood of the desired one (and determining quantitatively the size of this neighborhood is a clear open issue). This is due to the fact that a redundant number of visual features, typically the coordinates of at least four image points, are used as inputs of the control scheme for avoiding the famous cylinder of singularities when only three image points are considered [73, 99], and also since there usually exist four different poses such that the image of three points is the same [44]. Global asymptotic stability seems to be out of reach since the existence of attractive local minima has been exhibited in [20]. It concerned the case of coplanar points with a desired configuration not parallel to the image plane, following the result given in [80] that pose ambiguity exists in that particular case. More precisely, if a unique

pose exists from the perspective projection of four points, two symmetric solutions exist in the non-parallel case with the para-perspective model, leading to two similar (but not exactly the same) images for the poses corresponding to these two solutions with the perspective model. Designing one solution as the desired pose and starting from a pose near the other solution, it was not surprising that the system converges toward the latter, a local minimum, and not toward the former corresponding to the global minimum and desired pose. Furthermore, attractive local minima have also been recently exhibited in [84] for the case of four non coplanar points by looking at the behavior of the system near singularities found in the case of four points.

Even though they are known to exist, the local minima of image-based visual servoing systems have not received a lot of attention in the literature, possibly because of the theoretical difficulty of characterizing them [20]. Their existence have only been exhibited for some points' configurations by simulation, and the number, or even the existence, of local minima of arbitrary configurations is still an open problem. Clearly, these local minima represent a serious issue in real-world applications, and being able to formally characterize their properties and, in particular, their regions of attraction would represent a big step forward in the usability of these simple and robust controllers in practice.

It is well known that some of the classical image-based visual servoing controllers presents a natural Lyapunov function, at least under the assumption that no measurement noise is present [28, 67]. This motivated us to deeply investigate how we can exploit Lyapunov theory in the quest of studying the stability properties of image-based visual servoing systems. Lyapunov theory hardly needs presentations, being, arguably, one of the most widely used tools to assess the stability of dynamical systems [52, 60], which is broadly applied for both linear and non-linear systems, even if uncertain or non-continuous [33]. Lyapunov theory is intrinsically local in its original formulation and, while it allows for the analytical assessment of global stability in general (the so-called *LaSalle interpretation* of Lyapunov theory), evaluating the exact region of attraction of an equilibrium is a hard problem, even when the dynamics are known. The region of attraction's estimation problem has been extensively studied in the literature [10] and we can find several well known approaches to the problem. Classically, we can estimate the region of attraction by studying the invariant sublevel sets of a given Lyapunov function. In the case of polynomial systems, the evaluation is performed solving optimization problems with linear matrix inequality constraints, exploiting the sum-of-squares relaxation in polynomial optimization [25, 26, 51]. It is possible to generalize this approach to

allow for parametric Lyapunov functions [25, 101], while other authors rely on Zubov’s method [104] to design a suitable Lyapunov function [88], even though, since in image-based visual servoing we have a natural Lyapunov function already, we will not explore these strategies.

Even though we have a substantial variation in the choice of method, all of the Lyapunov-based strategies that we found to estimate regions of attraction rely on the same general philosophy: we define a global optimization problem, whose solution let us identify the largest sublevel set of the Lyapunov function so that its Lie derivative is strictly negative inside the set. The literature is then mainly concerned about the best way to solve this optimization problem [25, 48]. However, the very heart of this philosophy causes two limitations with these approaches, which could be particularly severe in the context of image-based visual servoing. The first one is that the estimate can be overly conservative. Since we look for a sublevel set for which the Lie derivative is strictly negative, it is clear that we cannot go any higher than the smallest sublevel set which includes a second equilibrium. Thus, in image-based visual servoing, we surely cannot go any higher than the value of the Lyapunov function at the local minimum, which is known to be very similar to the one at the desired pose, at least when the points are coplanar [20]. The second issue is more technical, but still relevant in our investigation. While these approaches are methodologically sound, the global optimization problem can be very complicated to solve in practice, even when the considered dynamical system is defined in low-dimensional settings. The complex structure of the image-based visual servoing systems, coupled with their definition on $SE(3)$, makes us believe that applying these methods is hardly a viable path, and our early attempts, carried out on simplified visual servoing systems¹, turned out to be intractable.

A completely different approach comes from the mathematics literature on dynamical systems theory, in particular from that part of mathematics that stems from Morse theory [74, 78]. The original objectives of M. Morse were almost completely unrelated from the stability analysis of dynamical systems, given that the aim of Morse theory is to assess the topological properties of manifolds. However, the way in which these properties are assessed bear a striking resemblance with some of the ideas coming from Lyapunov theory: we define a function on the manifold of interest, called a *Morse function*, and we study a class of dynamical systems, called *gradient-like systems*, which are such that

1. We considered that the camera is only allowed to move on a plane and only rotate with respect to the axis perpendicular to the plane, restricting its state-space to (a subset of) $SE(2)$.




Figure 1 – A stylized torus (i.e., the function’s domain) and the corresponding Reeb graph (a) and connectivity tree (b) for the height function.

the Lie derivative of the Morse function along the gradient-like flows is strictly negative, outside of the Morse function’s critical points. The gradient-like systems also have some additional hypotheses on their behavior when close to the critical points, thanks to which it is possible to precisely characterize the topological properties of the underlying manifold by studying the Morse function’s critical points.

A well-known example that comes from Morse theory is the celebrated *Reeb theorem* [74, Theorem 4.1], that tells us that, if we have a differentiable function whose domain is a compact manifold, and such that it only has two nondegenerate critical points, then the manifold is *homeomorphic* (i.e., topologically equivalent) to a sphere. Another famous example, which will be useful in the next section, is the so-called *Reeb graph*, which is a graphical tool which reflects the evolution of the level sets of a function whose domain is a manifold. The idea behind the Reeb graph is quite straightforward: for each level set, we consider that all the points that lie in the same *connected component* of the level set are equivalent, and we simply “collapse” each component to a dot. By drawing the evolution of the level sets’ components as the function increases, we obtain a *graph*, whose nodes represent the function critical points.

We can see a classical example of a Reeb graph in Figure 1a, where we consider the




Figure 2 – The error function and the corresponding gradient flow, showing a two dimensional saddle of index 1 separating two regions of attraction.

evolution of the height function on a (stylized) torus. As we move along the z -axis, the function first has a single component (in the range $-3 < z < -2$), then, when it meets the first saddle point, it splits into two components (for $-2 < z < 2$) which finally get rejoined at the second saddle point's level ($z = 2$). We clearly see this evolution on the Reeb graph in the figure, and, intuitively, we can use this tool to detect *holes* in the manifold, even when it is high-dimensional and impossible to visualize.

Later authors, and S. Smale in particular, brought Morse theory and dynamical systems' theory closer together, with the development of the so-called *Morse-Smale theory* [8], that, under stricter hypotheses than Morse, paints a precise relationship between the stable and unstable manifolds of gradient-like systems and the topology of their state-space. Further developments in this field eventually moved closer towards the stability and the regions of attraction of nonlinear dynamical systems: in [27], H.-D. Chiang and M. W. Hirsch showed that, given a nonlinear system and a stable equilibrium point for it, the boundary of the equilibrium's region of attraction is the union of the stable manifolds of other, unstable equilibria, under some generic-enough hypotheses that we will discuss soon. This characterization gives a crucial importance to unstable equilibria in studying the stability of dynamical system, which will be a central topic in our work.

Even though it is very difficult to prove this result in its full generality, the idea

behind this property is quite intuitive, and we want to illustrate it with a simple example. Consider the function $v(\mathbf{x}) = x_1 - 4x_1^2 + x_1^4 + 2x_2^2$ and its corresponding gradient flow, i.e., the dynamical system defined by $\dot{\mathbf{x}} = -\nabla v(\mathbf{x})$, both depicted in Figure 2. The function has two minima which correspond to stable equilibria of the gradient flow, towards which almost all trajectories converge to. By looking at the vector field, we can see that the saddle point acts as a “separator” between these minima: trajectories are not allowed to cross its stable space (depicted in red), and so trajectories starting on either “side” of the stable space must converge to the minimum on the same side. We can show that the saddle’s stable space stretches towards infinity (in fact, it is easy to see that it is the x_2 -axis), meaning that $\mathbb{R}_+^2 = \{\mathbf{x} \in \mathbb{R}^2 : x_1 > 0\}$ and $\mathbb{R}_-^2 = \{\mathbf{x} \in \mathbb{R}^2 : x_1 < 0\}$ are the regions of attraction of the local and global minima, respectively, while the saddle’s stable space is the boundary between them. We have shown that the saddle’s stable space represents the boundary of both regions of attraction of this simple example, and Chiang & Hirsch confirm that this is always the case, as long as the equilibria on the boundary are hyperbolic, that all trajectories on the boundary converge towards one of the equilibrium points and that the system is *structurally stable*².

As powerful (and beautiful) as this approach is, its direct application in the image-based visual servoing context seems unrealistic. Verifying that the required hypotheses hold for our systems is a big challenge *per se*; even taking them for granted, we would need to compute all of the system’s equilibria (unstable ones included, which would be practically impossible to find via simulation) and their stable and unstable manifolds. However, throughout the thesis we will be able to see how these seemingly hopelessly abstract ideas influenced our analysis.

All in all, in developing our own methodology to study the stability and the regions of attraction of image-based visual servoing systems, we were able to take several elements from both worlds, which ultimately allowed us to fill a gap that has existed in the visual servoing literature for more than 25 years. As we show in the next section, our approach gives us a new point of view on how to attack these problems, and, we believe, it strikes a good balance between the information required from the system’s dynamics and its applicability to “real-world” problems.

2. Structural stability, and its connection with the so-called *transversality condition* [97], has been pioneered by J. Palis [82]. Interestingly, Smale was the doctoral advisor of both Palis and Hirsch.

Contributions

The contributions of this thesis are twofold. The first “branch” is mainly concerned with developing new methodological tools that allow us to study the stability and assess the regions of attraction of dynamical systems which satisfy a global Lyapunov-like condition. Basically all the important results concerning this first part of our work are discussed in Chapter 1.

The second branch of our contributions is more practical in nature, and it is mainly concerned with applying this newfound theoretical insight to study the stability and the regions of attraction of image-based visual servoing systems. We develop this second part of our work through Chapters 2, 3 and 4.

Finally, in Chapter 5, we present some preliminary works on the impact of uncertainties on the equilibria and stability of dynamical systems. Even though the final objective is the application of these strategies in the context of image-based visual servoing, we only focused on the methodological aspects in this chapter, testing our ideas on some simple, idealized dynamical systems.

In Chapter 1, we first introduce the *energy-decreasing* formalism, where we assume that we are given an *energy function* and that we want to study the stability properties of dynamical systems, that we call *energy-decreasing systems*, for which the Lie derivative of the energy function is strictly negative outside of the energy’s critical points. We can equivalently characterize them as those systems for which the energy along any non-constant trajectory is strictly decreasing, which further motives their name.

In Section 1.2, we show that the energy-decreasing systems are closely related to one another, and that their main stability properties can be directly inferred from the energy function alone. We prove that they all have the same equilibria (both stable and unstable), which coincide with the critical points of the energy function. Additionally, we show that the *type* of critical point (i.e., whether it is a minimum, a maximum or a saddle) uniquely determines the *stability* of the equilibrium (i.e., if they are sink, sources, etc.) for any energy-decreasing system. These properties are intrinsically local in their formulation, and our investigation is clearly influenced by Lyapunov-like arguments; this is perfectly showcased by Lemma 1.1 (Section 1.2.2), in which, as a first step to show that the type of equilibrium does not change, we provide a generalization of the classical Lyapunov stability theorem for LTI systems which holds for any kind of equilibrium, and not only for stable ones.




Figure 3 – Comparison between the state-of-the-art estimates and the sublevel set at the saddle point for the function shown in Figure 2, together with an energy-decreasing trajectory starting just outside the saddle’s level.

It is well known that the sublevel sets of a Lyapunov function are invariant sets, and it is clear that, for the characterization of energy-decreasing systems we have made, any sublevel set of the energy function will be invariant for all energy-decreasing systems. As we will see in Section 1.3, this allows us to talk about regions of attraction that are agnostic to the choice of the particular energy-decreasing system, which, with a little abuse of notation, we will call *controller-independent regions of attraction*.

As we have seen in the state-of-the-art section, there are several classical approaches that try to find the largest possible controller-independent region of attraction, but we argued that, especially in visual servoing, these approaches can yield very conservative results. Looking at the example shown in Figure 2, however, it is intuitive to think that increasing the energy level up to the one of the saddle point would improve the estimation, as long as we focus on the sublevel set’s *component* connected to the minimum instead of considering the entire sublevel set. We show a comparison between the state-of-the-art estimates and the sublevel set at the saddle point in Figure 3, where we show, in blue, the level set corresponding to the energy level at the local minimum, which is the best estimate for the current state-of-the-art. In fact, we will be able to show that the (connected component of the) saddle’s sublevel set represent the largest possible controller-independent region of attraction for the minimum, and we showcase this by plotting an energy-decreasing trajectory (in red in the figure) which, starting just outside of the region, converges to the other local minimum.

This idea motivated us to develop a tool with which we can study the energy function and, in particular, how its sublevel sets are interconnected, which we call the *connectivity tree*, which is the heart of our theoretical contributions on the study of energy-decreasing systems. The connectivity tree is a graphical tool that tells us, for each energy level, how many components the sublevel set has (thus, in particular, whether it is connected or not). Its definition is clearly similar to the Reeb graph's one, the difference being that we consider the components of the *sublevel* sets for the connectivity tree, while we consider the components of the *level* set for the Reeb graph. We can see a comparison between the Reeb graph and the connectivity tree for the height function on the torus in Figure 1. From a topological point of view, the connectivity tree carries less information than the Reeb graph, but it is much easier to construct: the Reeb graph construction algorithm usually requires an explicit description of the manifold under the form of a simplicial mesh, with close-to-linear Lyapunov function on each simplex (see, e.g, [85]), which is simply unrealistic in our context. By contrast, to construct the connectivity tree of an energy function, we just need to know all its critical points and be able to do a finite number of reliable simulations of the gradient flow (or, alternatively, any energy-decreasing system).

The most important contribution of the connectivity tree is that we can use it to find the *largest possible* controller-independent region of attraction for *each minimum* (local or global) of the energy function. It also allows us to discuss a weaker form of the region of attraction, within which we allow the presence of other equilibria, as long as they are unstable. We speak in this case of *probability-1 region of attraction*, whose name is due to the fact that, for almost any initial position in the region, the trajectories of any energy-decreasing system will converge to the minimum. The connectivity tree allows us to find regions of attraction that potentially have a significantly larger size than classical estimates, examples of which we will see in Section 1.3.

The connectivity tree is only applicable for those sublevel sets which are compact, but the image-based visual servoing energy function have at least some unbounded level sets. We were able to identify, in Section 1.4.1, two situations so that the sublevel sets of the energy function become non-compact. The first one is very well known in the field, and it is due to the presence of trajectories where the camera runs off to infinity [20]. The other one though is seemingly unknown in the field, and it is due to trajectories where the camera tries to converge towards the object, while decreasing the energy at all times. This is a surprising behavior, especially because we expect that the energy goes to infinity as the camera gets closer to the object, the reason why we decided to call this peculiar

behavior a *hole-in-the-wall*. Analyzing these situations, we are able to compute a critical energy level that guarantees compactness of all sublevel sets below it, which allows us to make sense of the connectivity tree.

We now turn to the investigation of the second, more practical branch of our contributions. In Chapter 2, we study the critical points of the energy function of a class of classical image-based visual servoing systems, namely the control schemes that use the Cartesian coordinate of N points as visual features, which are known to be energy decreasing. In Section 2.1, we propose a nonlinear change of variables that allows to formulate the equilibrium condition directly in the *extended features*' space. These new variables make the exhaustive and verified computation of all equilibria for these systems a tractable problem, allowing, for the first time ever in the literature, to find and study their unstable equilibria and to verify the number of local minima. We also show that, if we take the camera's state-space as domain, the change of variables is actually a bijection, allowing to recover the corresponding equilibria in the camera's state-space. This is discussed in Section 2.2, where we also provide an algorithm that allows to reconstruct the camera's pose.

The exhaustive computation of all equilibria, together with the identification of the critical energy level that guarantees compactness of all sublevel sets below it, allows us to compute the connectivity tree of the image-based visual servoing energy function. In Chapter 3, we apply our new methodology to a set of representative test cases. We first compute all the equilibria and the corresponding heteroclinic orbits, for which we use the Julia package `VisualServoingToolbox.jl`³, developed by the author. Then, using this information, we compute the connectivity tree and the controller-independent regions of attraction for all minima, shining a new light on the dynamical properties of these robotic systems.

Our change of variables allows for the equilibria computation of other well-know non-energy-decreasing visual servoing systems, like the ones using an approximated version of the interaction matrix in the control loop, for which we analyze a set of representative test cases in Chapter 4. These approximated controllers have some advantages that are well known in the field, and there have been several attempts in the literature to combine them with the energy-decreasing ones. This motivated us to propose our own way of combining them, making use of our new methodological tools.

This analysis led to the synthesis of the *meta-controller*, that is discussed in Sec-

3. <https://github.com/acolotti/VisualServoingToolbox.jl>

tion 4.3. The meta-controller gives rise to an energy-decreasing system, and the intuition behind it is that it wants to combine the “good behaved” trajectories of energy-decreasing systems with the uniqueness of the minimum of non-energy-decreasing ones. We will see, in Section 4.4, that this results in a substantial increase in size of the desired pose’s region of attraction.

Finally, in Chapter 5, we collected some preliminary results, only tested on some toy examples, on how to deal with uncertainties in the feedback loop. Throughout the thesis we systematically assume that the system’s state and velocities are always perfectly known, which is too idealistic of an hypothesis. We explored two approaches to handle uncertainties. The first one, discussed in Section 5.1, aims to characterize the impact of unmodeled uncertainties on the closed-loop system. The philosophy behind it is that, in practice, we can accept a small imprecision on the final position of the robot, as long as its macro-behavior is as expected. Based on this idea, we propose a new strategy that allows to quantitatively assess the impact of uncertainties by identifying the system’s *attractor* and its region of attraction. Talking about attractors generalize the concept of stable equilibria and allows to discuss the *practical stability* of the system even when the stable equilibria are not robust to perturbation.

The second part of the chapter is perhaps the least developed, and the farthest from the image-based visual servoing applications. We shift our attention to *sliding mode controllers*, which are known to be not sensitive to (small enough) unmodeled uncertainties. We propose, in Section 5.2, only a technical contribution which concerns the convergence properties of the so-called *super twisting algorithm*, where we provide new gain bounds that ensure convergence with time-varying perturbations.

Thesis’ structure

This manuscript is organized as follows.

Chapter 1 is dedicated to the properties of energy-decreasing systems. We first show that there is a link between these systems and a class of classical controllers in image-based visual servoing, which motivates our investigation (Section 1.1). Then, we flesh out all the properties that these systems have in common, like their fixed points and the *type* of the equilibria, as well as the minimal hypotheses that we need to ensure these properties (Section 1.2). We continue to what can be considered the heart of this chapter, in which we propose a new graphical tool, the *connectivity tree*, that allows us to easily study

the interconnection of the sublevel sets of an energy function. We also propose a simple algorithm to compute it, which, by extension, lets us discuss the *controller-independent* region of attraction of a given stable equilibrium (Section 1.3). We end the chapter by showing that, when we use the points' Cartesian coordinates as visual features, the natural energy function in image-based visual servoing has some non-compact sublevel sets, and characterize the maximal energy level that guarantees their compactness (Section 1.4).

In Chapter 2, we study a class of energy-decreasing image-based visual servoing controllers (namely, the control schemes that use the Cartesian coordinate of N points as visual features), for which we derive tractable polynomial models for their equilibrium condition. A first polynomial model is naturally expressed in the state-space of the system, but turns out to be intractable. We then propose a second model, expressed in an extended space of visual features, which lead to systems of equations that are medium size and strongly non-linear (Section 2.1). We then show that we can uniquely recover the equilibria expressed in camera state-space by using the solution of the *orthogonal Procrustes problem* (Section 2.2). We finally present the computational algebraic geometry methods that can solve these systems of equations (Section 2.3).

In Chapter 3, we apply our methodology to study the stability and the controller-independent regions of attraction for a benchmark of typical configurations of image-based visual servoing systems, where, due to computational constraints, we only consider configurations of $N = 4$ points. Thanks to a fruitful collaboration with Jorge García Fontán and Mohab Safey El Din (LIP6 - Sorbonne Université), we were able to use the `msolve` solver [9], which is a recent, highly optimized implementation based on Gröbner bases computations, to compute all fixed points. This exhaustive equilibria computation shows a striking regularity in the zoology of equilibrium points, while the resolution timings, ranging from fractions of a second for simple configurations and up to 35 hours for general ones on a 12 core server, demonstrate the difficulty of solving these problems. We then compute the connectivity tree for each example, and use `IBEX` [5, 17, 18, 102], which is a C++ library for constraint processing over real numbers based on interval arithmetic, as a verified system solver to find a guaranteed enclosure of the energy function's sublevel sets.

Then, in Chapter 4, we turn our attention to a class of non-energy-decreasing image-based visual servoing controllers. We show that the polynomial models derived in the previous chapter are still applicable in this case (Section 4.1), and we use once again `msolve` to compute all fixed points for a benchmark of typical configurations (Section 4.2). We then

show how we can exploit our newfound theoretical insight of energy-decreasing systems to propose a new composite controller, called *meta-controller*, that combine the large region of attraction of the non-energy-decreasing controllers with the nice convergence properties of the energy-decreasing ones (Section 4.3). We prove that the meta-controller is at least always continuous, and we display the improvement of the global minimum's region of attraction by applying it to image-based visual servoing (Section 4.4).

In Chapter 5 we can find some preliminary results on how to deal with uncertainties. The content of this chapter is still far from being applicable in image-based visual servoing, and, so, we focus mainly on the methodological aspects of our results, applying them to simple toy examples. We explored two approaches to handle uncertainties: the first one analyzes the so-called *practical stability* of the perturbed system, trying to identify a bound on the error that guarantees a good macro-behavior of the system (Section 5.1). The second approach aims at using a sliding mode controller, the *super-twisting algorithm*, which is known to reject perturbations. Our contribution here is very technical and far from the visual servoing applications, since we provide new gain bounds that ensure convergence with time-varying perturbations (Section 5.2).

EQUILIBRIA AND STABILITY OF ENERGY-DECREASING SYSTEMS

In this chapter, we want to argue why the exhaustive identification of all the equilibria for a well-known class of image-based visual servoing controllers deeply characterizes their behavior. We focus on *energy-decreasing* dynamical systems, which, intuitively, are systems with an associated *energy function* which decreases along all trajectories, since several classical visual servoing controllers are known to give rise to such a closed-loop system.

The chapter is organized as follows. In Section 1.1, we present three classical image-based visual servoing controllers that are known to be energy-decreasing in the visual servoing literature. We then prove, using a new approach based on the change of the Riemannian metric on $SE(3)$, that we can see these systems as *gradient systems*, and show how much information on their dynamics can be derived from the study of the energy function alone. Motivated by this, we shift our focus to arbitrary energy-decreasing dynamical systems, and, in Section 1.2, we generalize the properties derived in the context of gradient systems to this larger class of systems, in particular for what concerns the equivalence of their equilibria.

We will see that we put in a lot of effort to study the *unstable* equilibria of these systems, and the reason for this is presented in Section 1.3. In this section, we present one of the main contributions of the chapter, which is the construction of the *connectivity tree*, a graphical tool that allows us to easily identify the *controller-independent regions of attraction* of energy-decreasing systems.

Finally, in Section 1.4, after this long detour through the properties of energy-decreasing systems, we come back to the specifics of the energy function that naturally arises in image-based visual servoing. We analyze two peculiar behaviors of the classical image-based visual servoing controllers, and we use them to identify two critical values that are such that the corresponding *sublevel sets* of the energy function are non-compact, which

will have important consequences in the determination of the connectivity tree for these systems.

1.1 Gradient systems in image-based visual servoing

In this section, we first present a class of image-based visual servoing controllers that are known to be energy-decreasing, which we be one of the main *foci* of this thesis, and then discuss how these controllers share the equilibria and several associated properties.

To us, an energy-decreasing system is a time-invariant dynamical system with an associated scalar, state-dependent energy function such that it is strictly decreasing along all trajectories, except for the trivial ones where the system stays on an equilibrium. These can be seen as a generalization of *gradient systems* (see, e.g., [53]), which are dynamical systems whose vector field is defined as the negative gradient of a given energy function. Gradient systems are of interest because their dynamics are greatly simplified: for instance, the system's equilibria coincide with the energy function's critical points and there cannot be periodic or chaotic trajectories, among other interesting properties. We will see how this characterization fares for energy-decreasing systems in the next section.

In this thesis, we consider a traditional image-based visual servoing [23] task, where we take the image-plane coordinates of $N \geq 4^1$ points (not necessarily co-planar) as visual features. We assume that we have an ideal *pinhole camera*, that the points are visible from any direction and that the camera is controllable in velocity. The points' Cartesian coordinates ${}^o\mathbf{a}_i$, expressed in a given world frame, are constant and known *a priori* for the analysis. The visual features x_i, y_i of the i -th point ${}^c\mathbf{a}_i = (X_i, Y_i, Z_i)$, expressed in the camera frame, are:

$$x_i = \frac{X_i}{Z_i}, \quad y_i = \frac{Y_i}{Z_i} \quad (1.1)$$

With this choice of features, the interaction matrix \mathbf{L}_i related to $\mathbf{s}_i = (x_i, y_i)$ is given by [23]:

$$\mathbf{L}_i = \begin{bmatrix} -\frac{1}{Z_i} & 0 & \frac{x_i}{Z_i} & x_i y_i & -(1 + x_i^2) & y_i \\ 0 & -\frac{1}{Z_i} & \frac{y_i}{Z_i} & 1 + y_i^2 & -x_i y_i & -x_i \end{bmatrix} \quad (1.2)$$

Denoting with $\mathbf{s} = (\mathbf{s}_1, \dots, \mathbf{s}_N)$ the (stacked) image features and with $\mathbf{s}^* = (\mathbf{s}_1^*, \dots, \mathbf{s}_N^*)$

1. While the strategies proposed throughout the thesis also work for $N = 3$, we don't explicitly consider this case because it is such that the fixed points are located on the cylinder of singularities evoked in the Introduction.

the desired ones, the dynamics of the features error $\mathbf{e} = \mathbf{s} - \mathbf{s}^*$ are:

$$\dot{\mathbf{e}} = \mathbf{L}_e \boldsymbol{\tau}_c, \quad \mathbf{L}_e = [\mathbf{L}_1^\top, \dots, \mathbf{L}_N^\top]^\top, \quad (1.3)$$

where $\boldsymbol{\tau}_c = (\mathbf{v}_c, \boldsymbol{\omega}_c)$ is the camera's spatial velocity that represents the control inputs, with \mathbf{v}_c being the linear velocity and $\boldsymbol{\omega}_c$ being the angular one.

The general goal of an image-based visual servoing task is to make the image features \mathbf{s} converge to the desired features \mathbf{s}^* . Several strategies to design appropriate controllers for these tasks exist in the literature. We mainly focus in this work on three such strategies² [28, 67], that are all based on the general form:

$$\boldsymbol{\tau}_c = -\lambda \mathbf{C}_e \mathbf{e} \text{ with } \lambda \in \mathbb{R}_+. \quad (1.4)$$

More precisely, we consider:

- *Transpose* controller: $\mathbf{C}_e = \mathbf{L}_e^\top$,
- *Pseudo-inverse* controller: $\mathbf{C}_e = \mathbf{L}_e^+$,
- *Levenberg-Marquardt* controller: $\mathbf{C}_e = (\mathbf{L}_e^\top \mathbf{L}_e + \mu \mathbf{I})^{-1} \mathbf{L}_e^\top$, with $\mu \in \mathbb{R}_+$.

It is well known [28, 67] that these controllers give rise to energy-decreasing systems³. Their energy function $v(\mathbf{x})$ is defined as:

$$v(\mathbf{x}) = \frac{1}{2} \|\mathbf{s}(\mathbf{x}) - \mathbf{s}^*\|^2, \quad (1.5)$$

where $\mathbf{x} \in \text{SE}(3)$ represents the camera state, $\mathbf{s} \in \mathbb{R}^m$ the output vector (or features) and \mathbf{s}^* the desired features.

In fact, in [67], E. Malis goes a bit further in the characterization of these three controllers in discrete time, making an explicit link between control and optimization. In particular, it is shown that the transpose controller give rise to a *gradient-descent* algorithm, while the pseudo-inverse and Levenberg-Marquardt controllers give rise to a *Gauss-Newton* and, as the name implies, *Levenberg-Marquardt* algorithms, respectively (see, e.g., [79] for an overview on these optimization algorithms).

2. While these controllers will be the protagonists of our work, we will later consider, in Chapter 4, an *approximation* of these controllers where we assume that the depths Z_i are unknown, since they typically cannot be measured in practical applications.

3. In fact, the characterization of these controllers as energy-decreasing systems do not depend on the fact that we track the points' Cartesian coordinates. Thus, their analysis, that we will develop throughout the whole chapter (with the exception of Section 1.4), does actually hold even for different choices of visual features.

We are interested in these three controllers because, intuitively, since we can see them as optimization algorithms, we expect to be able to study them by analyzing the energy function alone. In fact, our aim is to make this intuition rigorous, and we will show in the following that it is possible to give a much stronger characterization to these systems: we will prove that they all give rise to *gradient systems*, i.e., if we look at the closed-loop system, their vector fields are the *gradient flow* (i.e., a dynamical system of the form $\dot{x} = -\nabla v(\mathbf{x})$) of the energy function.

This claim could appear quite surprising, since, in the classical intuition that we have from calculus, the gradient of a continuously differentiable function (and, by extension, its gradient flow) is unique, while the vector fields associated to these controllers are clearly different in general. However, this intuition relies on the rigid structure of Euclidean space.

In order to appreciate the difference with the case where our domain is a arbitrary smooth manifold, we want to briefly show how we usually define the gradient in Euclidean space, and what it takes to generalize it to functions defined on manifolds. The classical way of presenting the gradient of a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ goes as follows (see, e.g., [98]): we say that f is differentiable at \mathbf{x} if there is a linear transformation $\Lambda : \mathbb{R}^n \rightarrow \mathbb{R}$ such that:

$$\lim_{\mathbf{h} \rightarrow \mathbf{0}} \frac{\|f(\mathbf{x} + \mathbf{h}) - f(\mathbf{x}) - \Lambda(\mathbf{h})\|}{\|\mathbf{h}\|} = 0. \quad (1.6)$$

It is then possible to prove that, if f is differentiable, this linear transformation is unique and the (row) vector associated with it is $\nabla f(\mathbf{x})$.

This formulation certainly makes no sense when \mathbf{x} is an element of a manifold, since there is no such thing as “adding” together elements of a manifold. However, there is one well-known property that links the gradient and the directional derivative of f which can still make sense even in this situation. To define the directional derivative of f at \mathbf{x} in the direction \mathbf{d} , we can consider a smooth curve $\gamma(t) : (-\epsilon, \epsilon) \rightarrow \mathbb{R}^n$ such that $\gamma(0) = \mathbf{x}$ and $\dot{\gamma}(0) = \mathbf{d}$, and take the derivative with respect to t of $f(\gamma(t))$, computed at $t = 0$. Thus, by applying the chain rule, we get an explicit link between the gradient and the directional derivative of f as:

$$\left. \frac{\partial}{\partial t} f(\gamma(t)) \right|_{t=0} = \left(\nabla f(\gamma(t))^\top \dot{\gamma}(t) \right) \Big|_{t=0} = \langle \nabla f(\mathbf{x}), \mathbf{d} \rangle \quad (1.7)$$

where $\langle \cdot, \cdot \rangle$ is the standard inner product on \mathbb{R}^n .

It is possible to show that, in Euclidean space, this property is actually equivalent to

the definition of gradient that we discussed above. Thus, since all the operations involved in (1.7) can be carried out on (smooth) manifolds as well, it is possible to take it as the *definition* of gradient for a function whose domain is a manifold. Before formalizing this idea, though, we need to point out that both the direction \mathbf{d} and the inner product are now defined on the *tangent space* of the manifold at \mathbf{x} . Contrarily to the Euclidean space case, there is no natural choice of inner product on the tangent space, and, since each point of the manifold has its own tangent space, the inner product can change as we move on the manifold. We call the collection of inner products on a smooth manifold its *Riemannian metric*, and we will soon see that it plays a crucial role in the definition of gradient itself.

We formalize these ideas with the following definition, which is classical in the field of differential geometry [8].

Definition 1.1. *Let $v : \mathcal{M} \rightarrow \mathbb{R}$ be a smooth function on a manifold \mathcal{M} endowed with a Riemannian metric, i.e., for each $\mathbf{x} \in \mathcal{M}$ the tangent space $T_{\mathbf{x}}\mathcal{M}$ is endowed with an inner product $\langle \cdot, \cdot \rangle_{\mathbf{x}}$. We call the gradient of v with respect to the metric on \mathcal{M} the unique vector $\nabla v(\mathbf{x}) \in T_{\mathbf{x}}\mathcal{M}$ such that:*

$$D_{\mathbf{x}}v(\boldsymbol{\tau}) = \langle \nabla v(\mathbf{x}), \boldsymbol{\tau} \rangle_{\mathbf{x}}, \quad (1.8)$$

for all $\boldsymbol{\tau} \in T_{\mathbf{x}}\mathcal{M}$, where $D_{\mathbf{x}}v(\boldsymbol{\tau})$ denotes the directional derivative of v with respect to $\boldsymbol{\tau}$.

It should now be clear why we stressed the dependence on the Riemannian metric: in these more general settings, the gradient of a function is still uniquely determined, but only *up to a change in metric*. Since there is no natural choice of metric, it is indeed possible to define a whole family of gradient flows, which are all equivalent, in a sense which will be discussed later in the section.

We can see how this impacts the image-based visual servoing case by computing the gradient of the energy function (1.5). Taking the derivative with respect to time of $v(\mathbf{x}(t))$, and assuming that the pose and the velocity at zero are $\mathbf{x}(0) = \mathbf{x}_0$ and $\dot{\mathbf{x}}(0) = \boldsymbol{\tau}_c$, we get that the directional derivative of v in the direction $\boldsymbol{\tau}_c$ is:

$$\begin{aligned} \left. \frac{d}{dt} v(\mathbf{s}(\mathbf{x}(t))) \right|_{t=0} &= \left. \frac{\partial}{\partial \mathbf{s}} v(\mathbf{s}) \right|_{\mathbf{s}=\mathbf{s}(\mathbf{x}(0))} \left. \frac{\partial}{\partial \mathbf{x}} \mathbf{s}(\mathbf{x}) \right|_{\mathbf{x}=\mathbf{x}(0)} \left. \frac{d}{dt} \mathbf{x}(t) \right|_{t=0} = \\ &= (\mathbf{s}(\mathbf{x}_0) - \mathbf{s}^*)^\top \frac{\partial}{\partial \mathbf{x}} \mathbf{s}(\mathbf{x}_0) \boldsymbol{\tau}_c. \end{aligned} \quad (1.9)$$

Now, by the error's dynamics (1.3), we have that $\dot{\mathbf{s}} = \frac{\partial}{\partial \mathbf{x}} \mathbf{s}(\mathbf{x}) \boldsymbol{\tau}_c = \mathbf{L}_e \boldsymbol{\tau}_c$, obtaining:

$$\left. \frac{d}{dt} v(\mathbf{x}(t)) \right|_{t=0} = (\mathbf{s}(\mathbf{x}) - \mathbf{s}^*)^\top \mathbf{L}_e(\mathbf{x}_0) \boldsymbol{\tau}_c, \quad (1.10)$$

where we stress the dependency of the interaction matrix on the camera pose. Comparing (1.10) with the gradient definition in Definition 1.1, we get that $\nabla v(\mathbf{x})$ is the unique vector that satisfies:

$$\langle \nabla v(\mathbf{x}), \boldsymbol{\tau}_c \rangle_{\mathbf{x}} = (\mathbf{s}(\mathbf{x}) - \mathbf{s}^*)^\top \mathbf{L}_e \boldsymbol{\tau}_c, \quad (1.11)$$

which, considering the conventional inner product, clearly implies that:

$$\nabla v(\mathbf{x}) = \mathbf{L}_e^\top (\mathbf{s}(\mathbf{x}) - \mathbf{s}^*). \quad (1.12)$$

This derivation of the energy gradient is classical in the visual servoing literature, and it is the reason why we conclude that the transpose controller give rise to a gradient-descent algorithm, as discussed above (see, e.g., [28, 67]). However, this is not the only possible choice of inner product, and framing the change of controller as a *change of metric* on the manifold allows us to tell much more on their behavior.

Talking about a change of metric sounds like a hopelessly abstract argument. In practice, however, since any positive-definite bilinear form is an inner product, a change of metric boils down to looking for a symmetric positive-definite matrix \mathbf{P} such that $\langle \mathbf{u}, \mathbf{v} \rangle = \mathbf{u}^\top \mathbf{P} \mathbf{v}$. In our context, we can easily define some new metrics of interest (i.e., finding some symmetric positive-definite matrices) that, as we will see, allow us frame the three controllers introduced above as gradient systems.

We notice that the matrix $\mathbf{L}_e^\top \mathbf{L}_e$ is positive semi-definite, meaning that $(\mathbf{L}_e^\top \mathbf{L}_e + \mu \mathbf{I})$ is positive-definite, for any choice of $\mu \in \mathbb{R}_+$. Additionally, if we assume that \mathbf{L}_e is full column rank (i.e., \mathbf{x} is not a *singularity*), $\mathbf{L}_e^\top \mathbf{L}_e$ is positive-definite as well. Thus, we can use them to easily define new metrics on the camera state-space (as long as, for the last metric, \mathbf{x} is not a singularity), which are:

$$\begin{aligned} \langle \mathbf{u}, \mathbf{v} \rangle_{\text{SD}} &= \mathbf{u}^\top \mathbf{v} \\ \langle \mathbf{u}, \mathbf{v} \rangle_{\text{LM}, \mathbf{x}} &= \mathbf{u}^\top (\mathbf{L}_e(\mathbf{x})^\top \mathbf{L}_e(\mathbf{x}) + \mu \mathbf{I}) \mathbf{v}, \\ \langle \mathbf{u}, \mathbf{v} \rangle_{\text{GN}, \mathbf{x}} &= \mathbf{u}^\top (\mathbf{L}_e(\mathbf{x})^\top \mathbf{L}_e(\mathbf{x})) \mathbf{v} \end{aligned} \quad (1.13)$$

where we stress the dependency of \mathbf{L}_e on the camera state in order to highlight that the

inner products *change* as the camera moves, or, said differently, to highlight that, contrarily to the conventional inner product $\langle \cdot, \cdot \rangle_{\text{SD}}$, the metric has a non-trivial dependency on the tangent space itself.

We argue that defining the gradient as a function of the metric gives us a more abstract but stronger point of view on the considered controllers, since all the nice properties of gradient systems discussed at the beginning of the section are preserved for any choice of metric. Coupled with the simplicity of identifying the inner products with a positive-definite matrix, this allows us to easily show that the relationship between the three considered controllers runs much deeper than what appears on the surface. We have the following result.

Proposition 1.1. *The transpose and Levenberg-Marquardt controllers give rise to gradient systems for the energy function (1.5). Under the assumption that no singularities are present in the state space, the pseudo-inverse controller does as well.*

Proof. As discussed above, in order to prove the proposition it is sufficient to find suitable inner products so that we have the following equality:

$$\langle \mathbf{C}_e \mathbf{e}, \boldsymbol{\tau}_c \rangle_{\mathbf{x}} = (\mathbf{s}(\mathbf{x}) - \mathbf{s}^*)^\top \mathbf{L}_e \boldsymbol{\tau}_c, \quad (1.14)$$

with \mathbf{C}_e being either the transpose, pseudo-inverse or Levenberg-Marquardt version of \mathbf{L}_e . As we saw, the conventional inner product gives us $\nabla v(\mathbf{x}) = \mathbf{L}_e^\top (\mathbf{s}(\mathbf{x}) - \mathbf{s}^*)$. Consider now the inner product $\langle \cdot, \cdot \rangle_{\text{LM}, \mathbf{x}}$ defined in (1.13). We have:

$$\begin{aligned} (\mathbf{s}(\mathbf{x}) - \mathbf{s}^*)^\top \mathbf{L}_e \boldsymbol{\tau}_c &= (\mathbf{s}(\mathbf{x}) - \mathbf{s}^*)^\top \mathbf{L}_e (\mathbf{L}_e^\top \mathbf{L}_e + \mu \mathbf{I})^{-1} (\mathbf{L}_e^\top \mathbf{L}_e + \mu \mathbf{I}) \boldsymbol{\tau}_c = \\ &= \langle (\mathbf{L}_e^\top \mathbf{L}_e + \mu \mathbf{I})^{-1} \mathbf{L}_e^\top (\mathbf{s}(\mathbf{x}) - \mathbf{s}^*), \boldsymbol{\tau}_c \rangle_{\text{LM}, \mathbf{x}}, \end{aligned} \quad (1.15)$$

meaning that the Levenberg-Marquardt controller is a gradient system under the Riemannian metric induced by $\langle \cdot, \cdot \rangle_{\text{LM}, \mathbf{x}}$. Analogously, if we consider $\langle \cdot, \cdot \rangle_{\text{GN}, \mathbf{x}}$, with the same argument we can conclude that the pseudo-inverse controller is indeed a gradient system, under the assumption that \mathbf{x} is not a singularity. \square

The assumption that there are no singularities in the pseudo-inverse controller case is unfortunate but expected. In fact, it is well known that the Gauss-Newton algorithm is not, in general, error-decreasing at the points where the Jacobian matrix is singular (see [79]). Since the pseudo-inverse direction might increase the energy, it cannot represent the negative gradient direction, regardless of the metric.

This result shows that the equilibria for the three controllers are the same⁴, with the possible exception of the pseudo-inverse controller if a critical point of the energy function is on a singularity of \mathbf{L}_e . However, a simple algebraic argument comes to our rescue in this exceptional case: it is well known that the Moore-Penrose pseudo-inverse of a matrix has the same kernel of its transpose (see, e.g., [13, Theorem 1.2.2]). Thus, $\ker \mathbf{L}_e^+ = \ker \mathbf{L}_e^\top$, meaning that the pseudo-inverse controller does indeed have the same equilibria of the other two, regardless of the singularities.

The similarities between these three controllers do not end here. Their characterization as gradient systems allows us to discuss the *type* of equilibrium in an unified fashion. In general, we can characterize (isolated) equilibria of dynamical systems with their *index*⁵, i.e., the number of eigenvalues with positive real part of the vector field's Jacobian evaluated at the equilibrium. For example, an index-0 equilibrium is stable, while an index-1 one is a saddle with one unstable direction. For gradient systems, an equilibrium's index is completely determined by the spectrum of the energy function's Hessian matrix, computed at the equilibrium.

However, analogously to the gradient, the computation of the Hessian matrix depends on the chosen metric, so it is fair to ask whether the index can change depending on the metric. What we can show is that the equilibria's indexes are an intrinsic property of the energy function, regardless of the metric. Intuitively, the inner product can deform the vector space (i.e., it can change the *magnitude* of the eigenvalues), but, because of its positive-definiteness, it cannot change their signs. We formalize this intuition in the following result.

Theorem 1.1. *The index of a given equilibrium of a gradient system is preserved under change of metric.*

This result is definitely not new in the literature, but it is quite difficult to pinpoint an exact reference for it, and for this reason we decided to provide an original proof in Appendix A. In [95], Douglas Shafer hints at the fact that even stronger properties of gradient systems are considered to be folklore in the community, stating that:

It is well known that when the singularity at 0 is nondegenerate the topological type of $X = \text{grad}_g V$ is independent of [the metric] g , and that in fact vector-fields arising from different Riemannian metrics are not only [topologically] equivalent but conjugate in a neighborhood of 0.

4. By looking at the gradient definition (1.8), it is easy to convince ourselves that the gradient is zero if and only if all directional derivatives are null, which implies that $\nabla v(\mathbf{x}) = 0$ regardless of the metric.

5. Not to be mistaken with the *topological index*, which will be used in the following section.

Saying that two vector fields are topologically equivalent means that, in a neighborhood of the equilibrium, there exists a homeomorphism carrying oriented orbits of one vector field onto oriented orbits of the other, and it is possible to show that topological equivalence implies that the index of hyperbolic equilibria is the same for the two vector fields⁶.

Nevertheless, we not only believe that providing a self-contained proof of this fact could be interesting *per se*, but it actually provides motivation for looking into general energy-decreasing system. In fact, in the proof of Theorem 1.1 we proved something more⁷: the Jacobian matrix of a gradient vector field (i.e., the Hessian matrix of v under a given metric) might not be symmetric, but it is guaranteed to have real eigenvalues. We then have the following corollary.

Corollary 1.1. *The Jacobian matrix $\mathbf{J}_{\nabla v}(\mathbf{x})$ computed at \mathbf{x} has real eigenvalues, regardless of the metric.*

This result is important because, before rushing into the study of general energy-decreasing system, a legitimate doubt might arise: are we sure there are energy-decreasing systems which cannot be expressed as gradient systems, after a suitable change of metric?

To verify this, consider the error function $v(\mathbf{x}) = (1 + x_1)^2 + x_2^2 + (-2 + x_1^2 + x_2^2)^2$, and let $f(\mathbf{x})$ be defined as:

$$f(\mathbf{x}) = - \begin{bmatrix} \frac{\sqrt{2}}{2} & -\frac{\sqrt{2}}{2} \\ \frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2} \end{bmatrix} \nabla v(\mathbf{x}) = -\mathbf{R}\nabla v(\mathbf{x}). \quad (1.16)$$

It is easy to show that the system $\dot{\mathbf{x}} = f(\mathbf{x})$ is error decreasing with respect to $v(\mathbf{x})$. In fact, if we notice that $(\mathbf{R} + \mathbf{R}^\top) = \sqrt{2}\mathbf{I}$, we have that, for $\nabla v(\mathbf{x}) \neq 0$:

$$\dot{v}(\mathbf{x}) = -\nabla v(\mathbf{x})^\top \mathbf{R} \nabla v(\mathbf{x}) = -\frac{1}{2} \nabla v(\mathbf{x})^\top (\mathbf{R} + \mathbf{R}^\top) \nabla v(\mathbf{x}) = -\frac{\sqrt{2}}{2} \|\nabla v(\mathbf{x})\|^2 < 0, \quad (1.17)$$

where, in the second equality, we exploit the fact that the quadratic form associated to a given matrix is equivalent to the quadratic form associated to the matrix's symmetric part. The graph of $v(\mathbf{x})$ and the phase portrait of $\dot{\mathbf{x}} = f(\mathbf{x})$ are shown in Figure 1.1. The error function only has one minimum, one saddle point and one maximum. At this local maximum of the energy function (whose coordinates are $\mathbf{x}^* = \left(\frac{\sqrt{3}-1}{2}, 0\right)$), the Jacobian of the (smooth) energy-decreasing vector field $f(\mathbf{x})$ has two complex-conjugate eigenvalues,

6. This is due to the fact that the stable and unstable manifolds' dimensions are preserved under homeomorphism.

7. See the remark at the end of Appendix A.

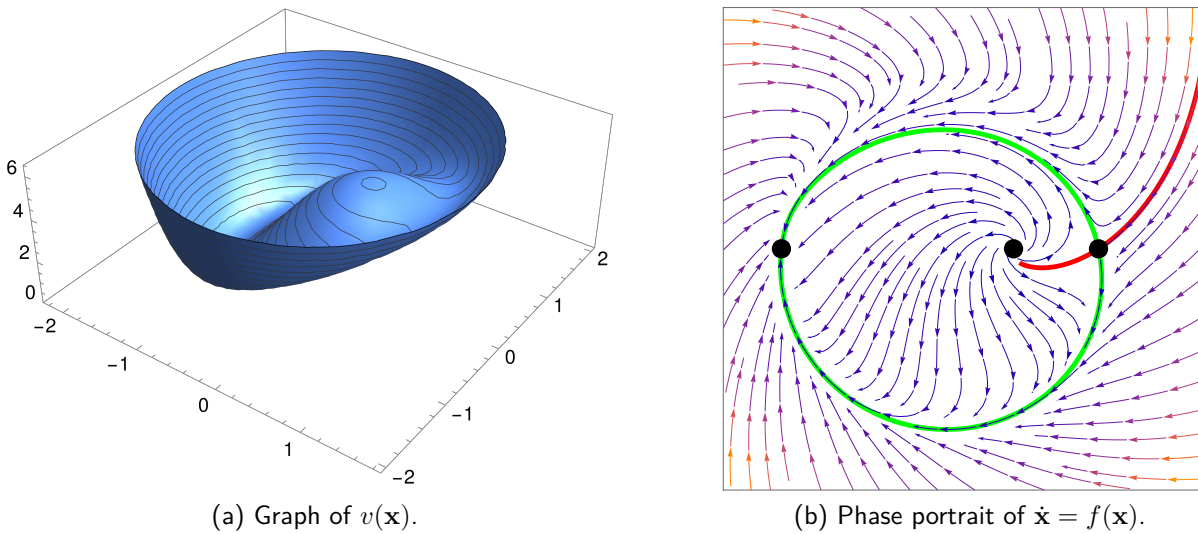


Figure 1.1 – (a): Graph of $v(\mathbf{x}) = (1 + x_1)^2 + x_2^2 + (-2 + x_1^2 + x_2^2)^2$. (b): Phase portrait of an energy-decreasing systems defined on $v(\mathbf{x})$. In green the unstable manifold of the saddle, in red its stable manifold.

and we can also see, in Figure 1.1b, that the trajectories are spiraling out of it. This is where Corollary 1.1 comes into play: since the Jacobian does not have real eigenvalues, this system cannot be a gradient system. With the doubt dispelled, and motivated by the fruitful analysis we developed so far, we indulge in the temptation of taking the generalization one step further, moving into the study of generic energy-decreasing systems.

1.2 Dynamical equivalence of generic energy-decreasing systems

In this section, we dig deeper into the link between energy-decreasing systems, weakening the hypotheses of only having gradient systems, i.e., systems whose vector field is the energy function's gradient under a suitable change of metric, and exploring which of their interesting properties are preserved in these settings. In Section 1.2.1, we first provide a more formal characterization of energy-decreasing system, and then show that all such systems have the same equilibria and that these equilibria represent the only points where the system can converge to. These properties hold true in great generality, even when the system is not differentiable at the critical points. In Section 1.2.2, we will need continuous differentiability to show that also the *type* of equilibrium is preserved.

While the non-differentiability hardly impacts the role of stable and repulsive equilibria, we will see two examples of non-differentiable energy-decreasing systems that showcase that saddle points can be fundamentally different in these cases.

We might think that this sounds like we are in for a lot of work just to give a strong characterization of the saddles, which seem to be of dubious practical interest. However, we will see in the next section how these unstable equilibria can play a fundamental role in the study of energy-decreasing systems, in particular for how they act as “separators” of the local minima’s regions of attraction, which will have interesting applications in visual servoing as well. The overall goal of these sections is to convince the reader that it is possible to infer important dynamical properties of any energy-decreasing system just by identifying the critical points of the associated energy function and studying how they are connected to each other.

From here on, we will always consider that an energy function is at least two times continuously differentiable and it only has *nondegenerate* critical points, i.e., the function’s Hessian computed at a critical point is always invertible. Additionally, unless stated otherwise, we consider that the energy function’s domain is an *n-dimensional smooth manifold*, i.e., a topological manifold (which, to us, is a space that is locally Euclidean of dimension n) together with a smooth *atlas* on it.

1.2.1 Fixed points

At the beginning of the chapter, we introduced energy-decreasing systems as dynamical systems with an associated energy function such that it is strictly decreasing along all trajectories, except for the trivial ones where the system stays on an equilibrium. This characterization takes into account whole trajectories, but we can equivalently characterize energy-decreasingness locally by saying that, at all points in the state-space outside of the energy’s critical points, the energy function’s directional derivative with respect to the vector field is strictly negative. We take this second point of view to provide our formal definition of energy-decreasing system.

Definition 1.2. *Let $v \in \mathcal{C}^2(\mathcal{M})$ and $f : \mathcal{M} \rightarrow T\mathcal{M}$ be locally Lipschitz⁸. We say that the system $\dot{\mathbf{x}} = f(\mathbf{x})$ is energy-decreasing for v if, for all $\mathbf{x} \in \mathcal{M}$, the Lie derivative is such that $\mathcal{L}_f v(\mathbf{x}) \leq 0$, with the equality holding if and only if \mathbf{x} is a critical point of v .*

8. For the results presented in the following, requiring the continuity of the vector field would actually be sufficient. However, local Lipschitzness entails existence and uniqueness of trajectories, which greatly simplifies the presentation.

First of all, a couple comments on this definition. Energy-decreasing systems clearly have a strong link with Lyapunov theory: if we take the classical Lyapunov stability theorem into account (see, e.g., [60]), we can immediately see, by restricting to a neighborhood of a local minimum of the energy function, that these local minima are stable equilibria for any energy-decreasing system. Analogously, by reversing time, we can conclude that local maxima of the energy function are repulsive equilibria.

However, Lyapunov theory does not help with saddle points, which, in a sense, represent one of the main protagonists of this section. As we will see, we need more sophisticated theoretical tools to prove that these critical points are indeed equilibria for any energy-decreasing system. Another area in which classical Lyapunov theory hardly helps is a *global* description of the system's behavior. In analogy with gradient systems, we are going to show that all non-diverging trajectories converge to a critical point of the energy function, as long as they are nondegenerate (which, again, will always be the case for us).

As promised, we have the following, comforting result.

Theorem 1.2. *Let $\dot{x} = f(x)$ be energy-decreasing for $v \in \mathcal{C}^2(\mathcal{M})$. Then, $x^* \in \mathcal{M}$ is an equilibrium for the system if and only if it is a critical point of v .*

Proof. We notice that, since the critical points of v are considered to be nondegenerate (and, thus, isolated), we can restrict the discussion to a neighborhood $U \ni x^*$ using local coordinates.

Let (U, φ) be a smooth coordinate chart on \mathcal{M} , with φ being a diffeomorphism from U to an open subset $\tilde{U} \subset \mathbb{R}^n$. Additionally, let $\hat{\varphi}$ denote the natural coordinates on the tangent bundle $T\mathcal{M}$. We can then define the local representations of v and f as $\tilde{v} = v \circ \varphi^{-1} : \tilde{U} \rightarrow \mathbb{R}$ and $\tilde{f} = \hat{\varphi} \circ f \circ \varphi^{-1} : \tilde{U} \rightarrow \mathbb{R}^n$, which let us discuss the properties of v and f in Euclidean settings.

In local coordinates, the Lie derivative $\mathcal{L}_f v(x)$ becomes simply $\nabla \tilde{v}(z)^\top \tilde{f}(z)$, with $z = \varphi(x) \in \mathbb{R}^n$. Thus, in order to prove the theorem, it suffices to show that $\tilde{f}(z) = 0$ if and only if $\nabla \tilde{v}(z) = 0$.

From the energy-decreasing definition, the \implies is trivial. For the other implication, we show that $-\nabla \tilde{v}$ and \tilde{f} have the same *topological degree* (see, e.g., [81]). Being $z^* = \varphi(x^*)$ isolated, we can find a neighborhood $\mathcal{N}(z^*)$ such that z^* is the only solution of $-\nabla \tilde{v}(z) = 0$ in $\mathcal{N}(z^*)$. Then, by [81, Proposition 3.2, Chap. IV], the degree $d(-\nabla \tilde{v}, \mathcal{N}(z^*), 0)$ is either 1 or -1 , depending on the sign of the Hessian's determinant at z^* . By defining the

(continuous) homotopy $H : [0, 1] \times \tilde{U} \rightarrow \mathbb{R}^n$ as:

$$H(t, z) = -t\nabla\tilde{v}(z) + (1-t)\tilde{f}(z), \quad (1.18)$$

we have that $d(\tilde{f}, \mathcal{N}(z^*), 0) = d(-\nabla\tilde{v}, \mathcal{N}(z^*), 0)$ if $H(t, z) \neq 0$ for all $t \in [0, 1]$ and $z \in \partial\mathcal{N}(z^*)$ (see [81, Proposition 2.4, Chap. IV]). For $t = 0, 1$ it is trivially true, while for all $t \in (0, 1)$ we have:

$$\nabla\tilde{v}(z)^\top H(t, z) = \underbrace{-t\|\nabla\tilde{v}(z)\|^2}_{<0} + \underbrace{(1-t)\nabla\tilde{v}(z)^\top \tilde{f}(z)}_{<0} < 0, \quad (1.19)$$

which implies that $d(\tilde{f}, \mathcal{N}(z^*), 0) \neq 0$ for any choice of $\mathcal{N}(z^*)$ sufficiently small. Thus, $\tilde{f}(z^*) = 0$, which concludes the proof. \square

We now discuss the global behavior of trajectories, for which, intuitively, we want to study where they converge to going forward or backwards in time. Formally, we do this by using the idea of α - and ω -limit sets, which allow us to drop the assumption of converging trajectories. Given a dynamical system $\dot{\mathbf{x}} = f(\mathbf{x})$, and denoting with $\mathbf{x}(t)$ an arbitrary trajectory, we call the ω -limit set of $\mathbf{x}(t)$ the set defined as [60]:

$$\omega(\mathbf{x}) = \left\{ \mathbf{p} \in \mathcal{M} : \exists \{t_n\} \rightarrow +\infty : \lim_{n \rightarrow +\infty} \mathbf{x}(t_n) = \mathbf{p} \right\}. \quad (1.20)$$

Analogously, the α -limit set of $\mathbf{x}(t)$ is defined as:

$$\alpha(\mathbf{x}) = \left\{ \mathbf{p} \in \mathcal{M} : \exists \{t_n\} \rightarrow -\infty : \lim_{n \rightarrow +\infty} \mathbf{x}(t_n) = \mathbf{p} \right\}. \quad (1.21)$$

The study of α - and ω -limit sets of gradient systems is classical in dynamical systems' theory (see, e.g., [53]). We adapt the following well known result to energy-decreasing systems, where we don't need the assumption that v only has nondegenerate critical points.

Proposition 1.2. *Let \mathbf{z} be an α - or ω -limit point of an energy-decreasing system's trajectory. Then, \mathbf{z} is an equilibrium.*

Proof. Assume that $\mathbf{z} \in \omega(\mathbf{x})$. We want to show that the energy is constant along the flow $\varphi(\mathbf{z}, s)$. Since $\mathbf{z} = \lim_{n \rightarrow +\infty} \mathbf{x}(t_n)$ for some sequence $\{t_n\} \rightarrow +\infty$, then, by the continuity of v and the fact that the energy decreases along trajectories, we have that $a = v(\mathbf{z})$ is




Figure 1.2 – Graph of the “Mexican hat” function.

the greatest lower bound of $\{v(\mathbf{x}(t)) : t \geq 0\}$. However, we can obtain the flow $\varphi(\mathbf{z}, s)$ by shifting $\mathbf{x}(t)$, i.e., if $\mathbf{x}(t_n) \rightarrow \mathbf{z}$, then $\mathbf{x}(t_n + s) \rightarrow \varphi(\mathbf{z}, s)$. The energy of $\varphi(\mathbf{z}, s)$ must still be the greatest lower bound of $\{v(\mathbf{x}(t)) : t \geq 0\}$, meaning that $v(\varphi(\mathbf{z}, s)) = a$ for all $s \in \mathbb{R}$. Thus, $\dot{v}(\mathbf{z}) = 0$, meaning that it is an equilibrium by Theorem 1.2. If $\mathbf{z} \in \alpha(\mathbf{x})$, we have that it is an ω -limit point for $\dot{\mathbf{x}} = -f(\mathbf{x})$, which is energy-decreasing for $-v$, meaning that $\dot{v}(\mathbf{z}) = 0$ as well. \square

If v only has isolated critical points, then all trajectories either tend to a critical point or they run off to infinity. This property allow us to guarantee that energy-decreasing systems always behave “nicely”, excluding the possibility of stable orbits or chaotic attractors. Moreover, if all the sub-level sets of the energy function are compact (for instance, if $\mathcal{M} = \mathbb{R}^n$ and the energy function is *radially unbounded*), then all the trajectories converge towards a critical point, and, if we additionally suppose that the manifold \mathcal{M} is compact, then all *orbits* (i.e., the full trajectory’s curve, with the time spanning from $-\infty$ to $+\infty$) connect two critical points.

As a final remark, one could think that, if the critical points are not isolated, then the non-diverging trajectories still converge to some equilibrium. This is not the case: surprisingly, there are examples of energy functions for which the corresponding gradient system has trajectories whose ω -limit set is not a point. A famous example is the so-called *Mexican hat* function by Curry [36], of which we can see the graph in Figure 1.2. For the gradient system related to this function, it is possible to show that there exists a trajectory for which its ω -limit set is the unit circle [82, Ex. 3, Sec. 1.1]. Intuitively,

this is the case because there are trajectories that “wrap around” the unit circle infinitely many times while converging towards it. Even though this example is quite pathological, it shows that we cannot escape the assumption of having nondegenerate critical points without renouncing having "well-behaved" trajectories⁹.

1.2.2 Stable and unstable manifolds

We want now to investigate what is the relation between the indexes of the energy function’s critical points and the index of the vector field’s Jacobian at those points. As it is the case for gradient systems, we expect that the equilibrium type is the same for all energy-decreasing systems. Pursuing this result, we first show, in the following lemma, that this is the case for *linear* energy-decreasing systems on \mathbb{R}^n , by proving that the systems’ stable and unstable subspaces all have the same dimensions. We can consider this lemma to be a weak generalization of the classical Lyapunov stability theorem for LTI systems, since it holds for any kind of equilibrium but it only provides a sufficient condition for the equilibrium’s type.

Lemma 1.1. *Let $\dot{\mathbf{x}} = \mathbf{A}\mathbf{x}$ be an n -dimensional linear system and $\mathbf{P} \in \mathbb{R}^{n \times n}$ be an invertible symmetric matrix. Then, if:*

$$\mathbf{A}^\top \mathbf{P} + \mathbf{P} \mathbf{A} \prec 0, \quad (1.22)$$

with \prec denoting negative definiteness, the stable and unstable subspaces of \mathbf{A} and $-\mathbf{P}$ have the same dimensions. In particular, \mathbf{A} doesn’t have any eigenvalue with null real part.

Proof. Let \mathcal{A}_s , \mathcal{A}_u and \mathcal{A}_c denote the stable, unstable and center subspaces of \mathbf{A} and let α_s , α_u and α_c be their dimensions. Define analogously the subspaces \mathcal{Q}_s and \mathcal{Q}_u with respect to $\mathbf{Q} = -\mathbf{P}$, with κ_s and κ_u their dimensions.

We want to show that $(\mathcal{A}_s \oplus \mathcal{A}_c) \cap \mathcal{Q}_u = \{0\}$, where \oplus denotes the direct sum. Let $\mathbf{x}_0 \in (\mathcal{A}_s \oplus \mathcal{A}_c)$ and let $\mathbf{x}(t)$ denote the system’s trajectory such that $\mathbf{x}(0) = \mathbf{x}_0$. Then, $\|\mathbf{x}(t)\|$ is bounded for all $t \geq 0$. Now, let $\mathbf{x}_0 \in \mathcal{Q}_u \setminus \{0\}$ and define $v(\mathbf{x}(t)) = \mathbf{x}(t)^\top \mathbf{Q} \mathbf{x}(t)$. We have that:

$$v(\mathbf{x}(t)) = v(\mathbf{x}_0) + \int_0^t \dot{v}(\mathbf{x}(t)) dt, \quad (1.23)$$

9. In fact, another condition that excludes this kind of behavior in \mathbb{R}^n would be to significantly restrict the class of energy functions: non-diverging trajectories of a gradient system are guaranteed to have an ω -limit set which is a point if the energy function is real analytic, but not if it is \mathcal{C}^∞ [1].

where $\dot{v}(\mathbf{x}(t)) = \mathbf{x}(t)^\top (\mathbf{A}^\top \mathbf{Q} + \mathbf{Q}\mathbf{A})\mathbf{x}(t) > 0$ denotes the time derivative of v . Suppose that $\int_0^t \dot{v}(\mathbf{x}(t))dt \rightarrow \bar{v}$ as $t \rightarrow +\infty$, for some finite value $\bar{v} \in \mathbb{R}_+$. Then, $\dot{v}(\mathbf{x}(t)) \rightarrow 0$, meaning that, given the positive definiteness of $\mathbf{A}^\top \mathbf{Q} + \mathbf{Q}\mathbf{A}$, $\mathbf{x}(t) \rightarrow 0$. However, this would mean that $v(\mathbf{x}(t)) \rightarrow 0$ as well, which cannot happen since it is bounded from below by $v(\mathbf{x}_0) > 0$. Thus, $v(\mathbf{x}(t)) \rightarrow +\infty$, which implies $\|\mathbf{x}(t)\| \rightarrow +\infty$, concluding that there exists no $\mathbf{x}_0 \in (\mathcal{A}_s \oplus \mathcal{A}_c) \cap \mathcal{Q}_u \setminus \{0\}$. Analogously, by taking $\tilde{\mathbf{A}} = -\mathbf{A}$ and $\tilde{\mathbf{Q}} = -\mathbf{Q}$, we can show that $(\tilde{\mathcal{A}}_s \oplus \tilde{\mathcal{A}}_c) \cap \tilde{\mathcal{Q}}_u = (\mathcal{A}_u \oplus \mathcal{A}_c) \cap \mathcal{Q}_s = \{0\}$.

Now, we recall that, given two linear subspaces S_1 and S_2 of \mathbb{R}^n such that $S_1 \cap S_2 = \{0\}$, then $\dim S_1 + \dim S_2 \leq n$. Paired with the results above, it gives us the following relationships:

$$\begin{cases} \alpha_s + \alpha_c \leq n - \kappa_u = \kappa_s \\ \alpha_u + \alpha_c \leq n - \kappa_s = \kappa_u \end{cases}, \quad (1.24)$$

where we used the fact that $\kappa_s + \kappa_u = n$. Taking the sum of both equations, we obtain $\alpha_s + \alpha_u + 2\alpha_c \leq n$, which tells us that $\alpha_c = 0$ since $\alpha_s + \alpha_u + \alpha_c = n$. This leads to $\alpha_u + \alpha_s = n$, which, together with (1.24), shows that $\alpha_s = \kappa_s$ and $\alpha_u = \kappa_u$, concluding the proof. \square

Being stability a local property, we can prove that the same holds for general nonlinear systems on manifolds by restricting on a neighborhood of the equilibrium and studying the system's linearization in order to exploit Lemma 1.1. We have the following result.

Theorem 1.3. *Let $\dot{\mathbf{x}} = f(\mathbf{x})$ be energy-decreasing for $v \in \mathcal{C}^2(\mathcal{M})$, with $f \in \mathcal{C}^r(\mathcal{M})$, $r \geq 1$, and let $\mathbf{x}^* \in \mathcal{M}$ be an equilibrium of f . Then, if \mathbf{x}^* is hyperbolic, its stable and unstable (sub)manifolds, denoted with $W^s(\mathbf{x}^*)$ and $W^u(\mathbf{x}^*)$, are of class \mathcal{C}^r and have the same dimension of the unstable and stable subspaces, respectively, of $H_v(\mathbf{x}^*)$, which denotes the Hessian of v at \mathbf{x}^* .*

Proof. Analogously to Theorem 1.2, we can restrict to a neighborhood $U \subset \mathcal{M}$ of \mathbf{x}^* and consider local representations $\tilde{v} : \tilde{U} \rightarrow \mathbb{R}$ and $\tilde{f} : \tilde{U} \rightarrow \mathbb{R}^n$ under a smooth coordinate chart (U, φ) , with $\tilde{U} = \varphi(U) \subset \mathbb{R}^n$. Additionally, being \mathbf{x}^* both an equilibrium of f and a critical point of v , the eigenvalues' signs of $Df_{\mathbf{x}^*}$ and $H_v(\mathbf{x}^*)$ are the same of $D\tilde{f}_{\mathbf{z}^*}$ and $H_{\tilde{v}}(\mathbf{z}^*)$, with $\mathbf{z}^* = \varphi(\mathbf{x}^*)$, regardless of the choice of φ (see, e.g., [82, Sec. 2.4] and [8, Sec. 3.1]). We finally assume, without loss of generality, that $\mathbf{z}^* = \mathbf{0}$ (see, e.g., [66]).

In order to apply Lemma 1.1, we want to study the energy-decreasing condition as we

simultaneously linearize \tilde{f} and $\nabla\tilde{v}$. Their Taylor expansions at the origin are given by:

$$\begin{aligned}\nabla\tilde{v}(\mathbf{z}) &= H_{\tilde{v}}(\mathbf{0})\mathbf{z} + \|\mathbf{z}\|\rho_v(\mathbf{z}) \\ \tilde{f}(\mathbf{z}) &= D\tilde{f}_0\mathbf{z} + \|\mathbf{z}\|\rho_f(\mathbf{z})\end{aligned}, \quad (1.25)$$

where the vectors $\rho_v(\mathbf{z})$, $\rho_f(\mathbf{z})$ are the so-called *remainders*, i.e., functions such that $\rho_v(\mathbf{z}) \rightarrow \mathbf{0}$ and $\rho_f(\mathbf{z}) \rightarrow \mathbf{0}$ as $\mathbf{z} \rightarrow \mathbf{0}$ (see, e.g., [4, Chapter 9]). It is evident that the sum of remainders is still a remainder; additionally, given a matrix $\mathbf{B} \in \mathbb{R}^{n \times n}$ and a remainder $\rho(\mathbf{z})$, any function of the form $\frac{\mathbf{z}^\top \mathbf{B} \rho(\mathbf{z})}{\|\mathbf{z}\|}$ is a remainder as well, since, by Cauchy-Schwarz inequality:

$$\limsup_{\mathbf{z} \rightarrow \mathbf{0}} \left| \frac{\mathbf{z}^\top \mathbf{B} \rho(\mathbf{z})}{\|\mathbf{z}\|} \right| \leq \lim_{\mathbf{z} \rightarrow \mathbf{0}} \frac{\|\mathbf{z}\| \|\mathbf{B} \rho(\mathbf{z})\|}{\|\mathbf{z}\|} = 0. \quad (1.26)$$

Thus, by denoting $\mathbf{A} := D\tilde{f}_0$ and $\mathbf{P} := H_{\tilde{v}}(\mathbf{0})$, we can rewrite the energy-decreasing condition, for all $\mathbf{z} \in \tilde{U}$, as follows:

$$\begin{aligned}\nabla\tilde{v}(\mathbf{z})^\top \tilde{f}(\mathbf{z}) &= (\mathbf{P}\mathbf{z} + \|\mathbf{z}\|\rho_v(\mathbf{z}))^\top (\mathbf{A}\mathbf{z} + \|\mathbf{z}\|\rho_f(\mathbf{z})) = \\ &= \mathbf{z}^\top \mathbf{P}\mathbf{A}\mathbf{z} + \|\mathbf{z}\|^2 \underbrace{\left(\frac{\mathbf{z}^\top \mathbf{P}\rho_f(\mathbf{z})}{\|\mathbf{z}\|} + \frac{\rho_v(\mathbf{z})^\top \mathbf{A}\mathbf{z}}{\|\mathbf{z}\|} + \rho_v(\mathbf{z})^\top \rho_f(\mathbf{z}) \right)}_{\rho(\mathbf{z})} = \\ &= \frac{1}{2} \mathbf{z}^\top (\mathbf{A}^\top \mathbf{P} + \mathbf{P}\mathbf{A}) \mathbf{z} + \|\mathbf{z}\|^2 \rho(\mathbf{z}) \leq 0,\end{aligned} \quad (1.27)$$

where $\rho(\mathbf{z})$ is a remainder. Now, let \mathbf{z}_0 be a normalized eigenvector of $(\mathbf{A}^\top \mathbf{P} + \mathbf{P}\mathbf{A})$ corresponding to its largest eigenvalue, denoted with λ_0 . By restricting (1.27) along the direction of \mathbf{z}_0 , we can always find $\epsilon > 0$ such that $\alpha\mathbf{z}_0 \in \tilde{U}$ for all $\alpha \in [-\epsilon, \epsilon]$, leading to:

$$l(\alpha) := \nabla\tilde{v}(\alpha\mathbf{z}_0)^\top \tilde{f}(\alpha\mathbf{z}_0) = \alpha^2 \left(\frac{\lambda_0}{2} + \rho(\alpha\mathbf{z}_0) \right) \leq 0, \quad (1.28)$$

which clearly implies:

$$\lim_{\alpha \rightarrow 0} \frac{l(\alpha)}{\alpha^2} = \frac{\lambda_0}{2} \leq 0. \quad (1.29)$$

Both \mathbf{A} and \mathbf{P} are invertible by hypothesis, since we assume that \mathbf{x}^* is hyperbolic for $f(\mathbf{x})$ and, as always, that it is a nondegenerate critical point for the energy function. Thus, $\lambda_0 \neq 0$, meaning that all the eigenvalues of $(\mathbf{A}^\top \mathbf{P} + \mathbf{P}\mathbf{A})$ are strictly negative, i.e.:

$$\mathbf{A}^\top \mathbf{P} + \mathbf{P}\mathbf{A} \prec 0. \quad (1.30)$$

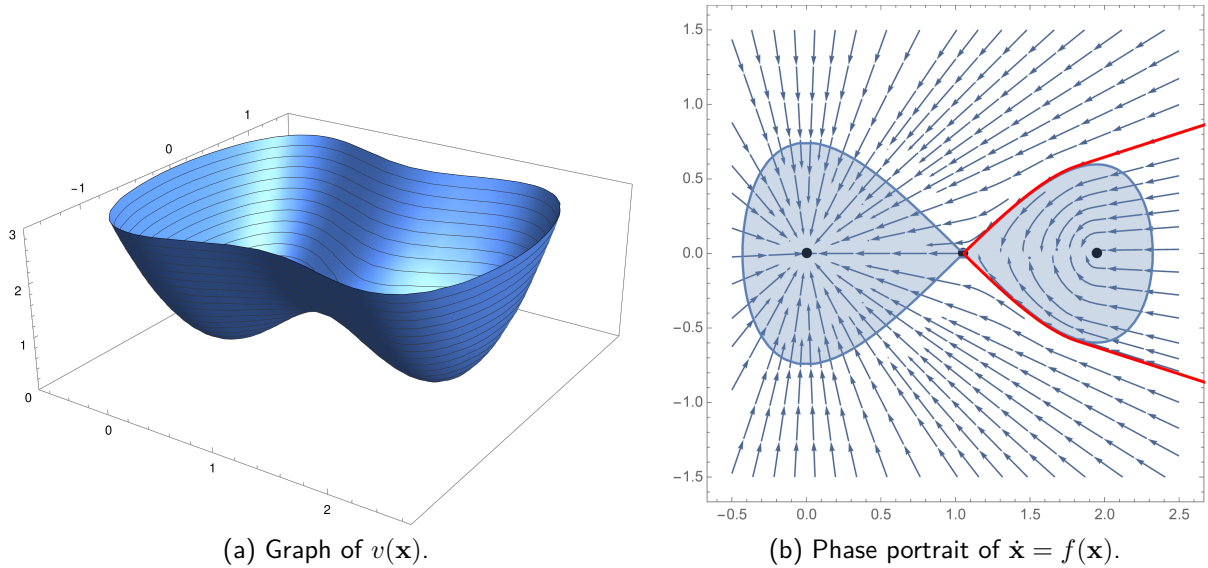


Figure 1.3 – Graph (left) of $v(\mathbf{x}) = 4.1x_1^2 - 4x_1^3 + x_1^4 + 2x_2^2$ and phase portrait (right) of the energy-decreasing system defined on $v(\mathbf{x})$. Red lines represent the stable manifold of the saddle, that separate the regions of attraction of the two minima.

We can then apply Lemma 1.1, concluding that the unstable subspace of \mathbf{A} have the same dimension of the stable subspace of \mathbf{P} , and vice versa.

Finally, since the eigenvalues' signs of $Df_{\mathbf{x}^*}$ and $\mathbf{A} = D\tilde{f}_0$ agree, $W^s(\mathbf{x}^*)$ and $W^u(\mathbf{x}^*)$ are \mathcal{C}^r submanifolds of \mathcal{M} with the same dimension of the stable and unstable subspaces of \mathbf{A} , respectively (see [82, Sec. 2.6], in particular Props. 6.1-6.2 and the section's conclusion). □

As discussed at the beginning of the section, even though Theorem 1.3 requires f to be (at least) continuously differentiable, we can easily relax this hypothesis when \mathbf{x}^* is a local minimum or a local maximum of v , since Lyapunov stability theorem guarantees that \mathbf{x}^* is an asymptotically stable or repulsive equilibrium, respectively, under the hypothesis that f is Lipschitz continuous. However, losing continuous differentiability of f can lead to unexpected results around saddle points, as the next examples show.

We consider an energy-decreasing control law defined using a strictly convex optimization problem that depend continuously on instantaneous data. This kind of construction is common in sensor-based control, and it will have important applications in Chapter 4. As we will see, such a command law is continuous under mild hypothesis, but easily fails to be differentiable: changes in active constraints produce loss of differentiability, and, even

though this does not usually happens on fixed points, we can easily design controllers that are not differentiable on fixed points.

Consider the energy function $v(\mathbf{x}) = 4.1x_1^2 - 4x_1^3 + x_1^4 + 2x_2^2$, whose global minimum is at the origin, but also has a local minimum and a saddle in between them. The graph of $v(\mathbf{x})$ is shown in Figure 1.3a. We first consider an auxiliary system that is not error decreasing but has good convergence properties toward the target, here for simplicity $f_0(\mathbf{x}) = -\mathbf{x}$, and then enforce error decrease by projecting it on a ball centered on $c = -R \nabla v(\mathbf{x})$ of radius $r = (R - \epsilon) \|\nabla v(\mathbf{x})\|$ (with $R = 100$ and $\epsilon = 0.01$) to obtain the following composite error-decreasing system:

$$f(\mathbf{x}) = \arg \min_{\mathbf{u} \in B(c, r)} \|\mathbf{u} - f_0(\mathbf{x})\|^2. \quad (1.31)$$

Note that when $\nabla v(\mathbf{x}) = 0$ we have $B(c, r) = \{0\}$ and hence $f(\mathbf{x}) = 0$. It is possible to show that this composite control law is continuous, even though we will introduce the right tools to prove it only in Chapter 4 (in fact, it is immediate to verify that the hypotheses of Proposition 4.1 are satisfied). The phase portrait of $\dot{\mathbf{x}} = f(\mathbf{x})$ is shown in Figure 1.3b, together with the stable manifold of the saddle in red, which separates the region of attraction of the two attracting fixed points. The interesting fact is that this stable manifold seems to be “broken”, which indicates a failure of Theorem 1.3 due to the non-differentiability of the vector field at the saddle.

However, the stable manifold still appears to be a 1-dimensional non-smooth manifold. Given this example, one could expect in this case that the system’s stable manifold still has the same dimension of the unstable subspace of the energy’s Hessian, while being non-smooth at the saddle point. Unfortunately, continuous differentiability of f is a crucial property without which the dynamics around the saddles can be radically different from the classical intuition provided by the stable manifold theorem, as the next example shows.

Consider the dynamical system $\dot{\mathbf{x}} = f(\mathbf{x})$ defined as:

$$f(\mathbf{x}) = \begin{cases} (-x_1, -x_2), & \text{if } |x_1| \leq |x_2| \\ (x_1 - 2|x_2|, -x_2), & \text{if } x_1 > |x_2| \\ (x_1 + 2|x_2|, -x_2), & \text{otherwise} \end{cases} \quad (1.32)$$

It is easy to verify that $f(\mathbf{x})$ is locally Lipschitz and that it is energy-decreasing with respect to $v(\mathbf{x}) = \frac{1}{2}(-x_1^2 + 2x_2^2)$, whose graph can be seen in Figure 1.4a. The energy function has only one critical point at $\mathbf{x} = (0, 0)$, where the Hessian’s spectrum is $\{1, -2\}$. However, the system’s phase portrait, which can be seen in Figure 1.4b, shows that the

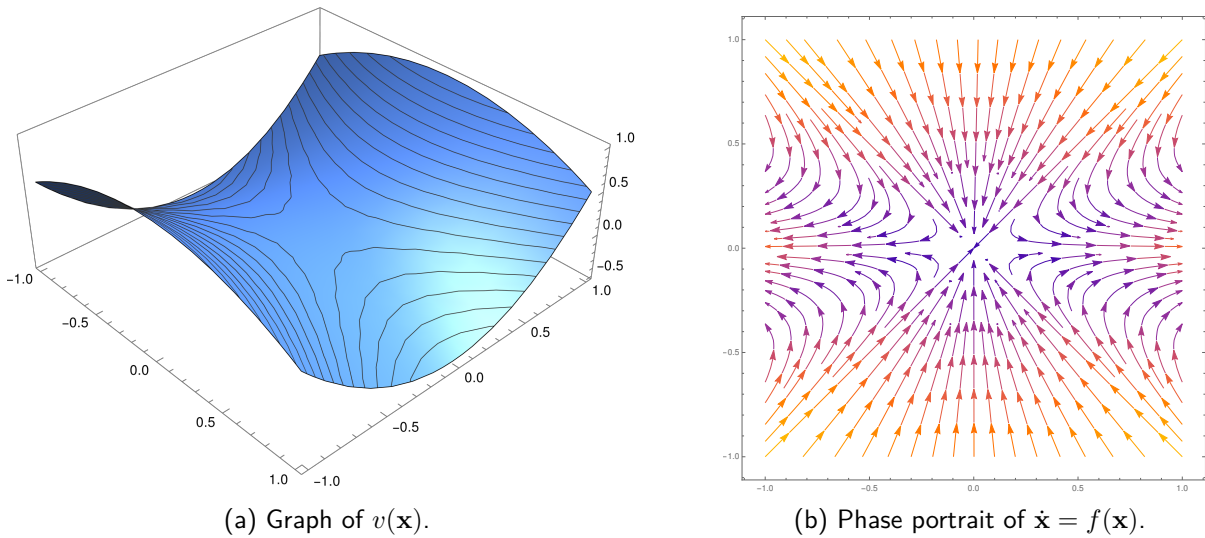


Figure 1.4 – Graph (left) and phase portrait (right) of $v(\mathbf{x}) = \frac{1}{2}(-x_1^2 + 2x_2^2)$ and the energy-decreasing system defined in (1.32).

equilibrium's stable set is not a manifold of dimension 1 (in fact, it appears to be a 2-dimensional manifold with boundary), definitely contradicting the conjecture that a non-smooth version of Theorem 1.3 might be true. Additionally, and quite surprisingly, it is possible to show that the unstable set is a 1-dimensional manifold (the x_1 -axis), meaning that not only the sum of their dimensions does not match the dimension of the state-space, but there does not seem to be any evident symmetry or pattern in how these peculiar sets split the space.

Together, these two examples seem to imply that the smoothness properties of f are indeed insufficient to characterize the system's behavior in an equilibrium's neighborhood when f is not at least continuously differentiable, and, most likely, it will be necessary to take additional properties into account to identify those cases where a non-smooth version of Theorem 1.3 might be applicable.

1.3 Importance of unstable equilibria: the connectivity tree

Up to now, we focused on fully characterizing the equilibria of energy-decreasing systems defined on some common energy function. Unearthing the deep relationship that exists between these systems is of self-evident theoretical interest; however, in more prac-

tical settings, the impact of unstable equilibria might be perceived as less relevant, since even small numerical errors are sufficient to destabilize trajectories converging towards saddles making them inevitably fall into local minima of the energy function.

In this section, we want to argue that saddle points play an important role even in practice, with key applications to robotics in general and image-based visual servoing in particular. We return to one of the examples provided in the Introduction, and use it as a stepping stone to show how saddle points have a crucial role in separating the regions of attraction of the energy function's local minima and, in particular, we will show that saddles with index 1 are the only unstable equilibria of interest in the assessment of how the regions of attraction are interconnected. We use the information coming from the knowledge of all critical points to propose a new theoretical tool, that we call *connectivity tree*, that allows us to visualize the interconnections of the energy function's sublevel sets, and we show how we can use this tool to propose the largest possible *controller-independent* region of attraction for any local minimum, i.e., the minimal set of initial states such that the corresponding trajectories are guaranteed to converge to the minimum for *any* energy-decreasing system. Additionally, we discuss a slightly weaker form of region of attraction, that we call *probability 1 region of attraction*, which can improve the region's size by introducing a relaxation that has no impact from a practical point of view: in this region, trajectories are guaranteed to converge for *almost any* initial condition in the set, i.e., there might be a zero-measure subset that converges to a different, unstable equilibrium.

1.3.1 State-of-the-art and its limits

The simple examples that we considered back in the Introduction hints at the fact that there is some sort of relationship between the boundary of regions of attraction and the stable manifolds of other, unstable equilibria. In particular, it seems reasonable to expect that these stable manifolds are (at least) part of the region of attraction's boundary, given that trajectories cannot cross them.

It turns out that not only this intuition is correct, but that it holds under very general conditions [27]: given a nonlinear system and a stable equilibrium point for it, it is possible to show that the boundary of the equilibrium's region of attraction is the union of the stable manifolds of other equilibria, as long as the system has only hyperbolic fixed points, it is *structurally stable*¹⁰ and that all trajectories on this boundary converge to an

10. Intuitively, a structurally stable system is such that the orbits connecting different equilibria do

equilibrium (notice that the system needs not to be energy-decreasing!).

However, while this characterization gives a rigid structure to the minima’s regions of attraction, it hardly provides a computationally-feasible approach for the identification of these regions. When their dimension is greater than one, the computation of stable manifolds in \mathbb{R}^n alone is an extremely difficult problem [61]. Additionally, we will later see that index 1 saddles play a special role in the separation of these regions, meaning that we would need to compute $(n - 1)$ -dimensional manifolds. In image-based visual servoing applications, we typically consider $n = 6$, making the computation simply unfeasible.

In the context of energy-decreasing systems, the high dimensionality is not the only limit of this approach: the region of attraction of a given minimum depends on the actual system, which for energy-decreasing systems defined on the same energy function. Given the unified point of view on energy-decreasing systems that we brought forth up until now, it makes sense to search for a strategy that allows us to focus exclusively on the energy function to propose at least an *approximation* of the minima’s regions of attraction that works for *all* energy-decreasing systems defined on the same energy function.

Once again, Lyapunov theory comes to save us. It is well known that the *sublevel sets* of a Lyapunov function are (positively) invariant sets for the corresponding dynamical system. This fact is widely used in the literature, in particular to provide *inner approximations* of the region of attraction: a typical approach is to “cut” the Lyapunov function at a value which is just below the level of the second-smallest equilibrium (see Figure 1.5a for an illustrative example), which is used, for instance, in [25, 48]. If this sublevel set is compact, then we can trivially apply LaSalle invariance principle (see, e.g., [60]) to show that, for any initial state taken in this set, all trajectories converge to the global minimum of the Lyapunov function.

If we now consider an energy function $v(\mathbf{x})$, the same construction would provide us with an inner approximation of the global minimum’s region of attraction that holds for *all* energy-decreasing systems for $v(\mathbf{x})$. It is then clear that this set is an inner approximation of the global minimum’s *controller-independent region of attraction*, which is defined, in the context of energy-decreasing systems, as the largest set for which these properties hold true for all energy-decreasing systems, or, equivalently, as the intersection of all the minimum’s regions of attraction for any energy-decreasing system.

not change under a small enough perturbation of the system’s parameters. Typical examples of *not* structurally stable systems are systems for which the unstable manifold of a saddle coincide with the stable manifold of a second one. For an introduction to structural stability, see [46, 82].

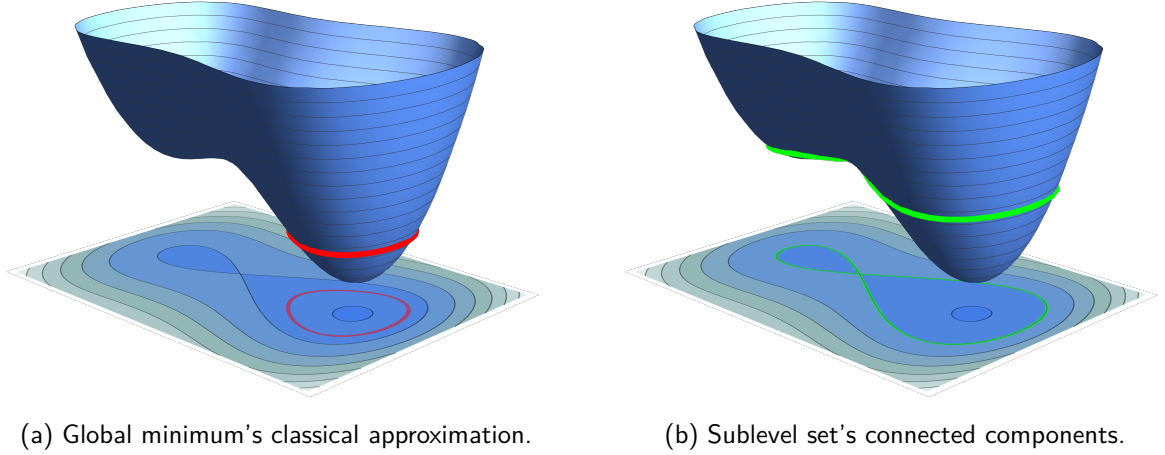


Figure 1.5 – Illustration of Lyapunov sublevel sets as inner approximations of the minima's regions of attraction.

1.3.2 Sublevel sets' connected components

We want to build upon this classical result to propose regions of attraction's estimates for any (possibly local) minimum, as well as improving the size of the estimate. A simple way to extend this result is to focus on the *connected component* of the sublevel set which includes the minimum of interest. Intuitively, if we have m local minima that are included in the sublevel set, and no other critical points, we expect that the sublevel set is formed by m disconnected component, each one representing an inner approximation of the corresponding minimum's controller-independent region of attraction. You can see in Figure 1.5 a simple example with two minima, for which it is evident that any sublevel set included between the green and red one has two disconnected components, which are then joined together once we reach the saddle point's energy level. We make this argument rigorous with the following result.

Theorem 1.4. *Consider an energy function $v \in \mathcal{C}^2(\mathcal{M})$ with isolated critical points, and let $\mathbf{x}^* \in \mathcal{M}$ be a local minimizer. Let \mathcal{V}_c be the sub-level set of v defined as:*

$$\mathcal{V}_c = \{\mathbf{x} \in \mathcal{M} : v(\mathbf{x}) \leq c\}, \quad (1.33)$$

and let \mathcal{V}_c^ be the connected component to \mathbf{x}^* of \mathcal{V}_c . Now, let c^+ be the minimum value of c such that $\mathcal{V}_{c^+}^*$ contains at least another critical point \mathbf{x}^+ , with $v(\mathbf{x}^+) = c^+$. Then, \mathbf{x}^+ is not a minimum and, for all $v(\mathbf{x}^*) \leq c < c^+$ such that \mathcal{V}_c^* is compact, \mathcal{V}_c^* is an inner*

approximation of the region of attraction of \mathbf{x}^ for any energy-decreasing system.*

Proof. We begin by noticing that, for all $c(\mathbf{x}^*) \leq c < c^+$, on the boundary of \mathcal{V}_c^* it holds $\mathcal{L}_f v(\mathbf{x}) < 0$, for any energy-decreasing system. Since it is compact, \mathcal{V}_c^* is then a (positively) invariant set for any such system [10]. We can thus apply LaSalle’s theorem [64], which says that the system converges to the largest (positively) invariant set in $\mathcal{Z}_c = \{\mathbf{x} \in \mathcal{V}_c^* : \mathcal{L}_f v(\mathbf{x}) = 0\}$, for all initial conditions within \mathcal{V}_c^* . Being $\mathcal{Z}_c = \{\mathbf{x}^*\}$ for all $v(\mathbf{x}^*) \leq c < c^+$, we have that any trajectory starting in \mathcal{V}_c^* converges to \mathbf{x}^* , meaning that \mathcal{V}_c^* is a subset of the controller-independent region of attraction of \mathbf{x}^* .

Finally, we prove by contradiction that \mathbf{x}^+ is not a minimum. If we assume that \mathbf{x}^+ is a minimum, then, being isolated, it’s possible to find a neighborhood $\mathcal{N}(\mathbf{x}^+)$ such that $v(\mathbf{x}) > v(\mathbf{x}^+)$, for all $\mathbf{x} \in \mathcal{N}(\mathbf{x}^+)$. In turn, by the definition of sub-level set, we have that $\mathcal{V}_{c^+}^* \cap \mathcal{N}(\mathbf{x}^+) = \{\mathbf{x}^+\}$, which means that \mathbf{x}^+ is an isolated point of the set $\mathcal{V}_{c^+}^*$. Clearly, this violates the connectedness of $\mathcal{V}_{c^+}^*$, thus proving the theorem. \square

Remark. *A sublevel set \mathcal{V}_c is always a closed set, since it is the preimage of the (closed) interval $(-\infty, c]$ through v , which is then closed by the continuity of v . If $\mathcal{M} = \mathbb{R}^n$, a sufficient condition for the compactness of \mathcal{V}_c is that it is bounded. However, when \mathcal{M} is an arbitrary smooth manifold, the compactness of its subsets is a significantly more subtle concept (see, e.g., [65]). We will see an example of a counter-intuitively non-compact subset in Sec. 1.4.2.*

Theorem 1.4 tells us that it is not possible to have two (or more) minima inside the same connected component without having at least another, unstable equilibrium that acts as a connection between them. This indeed confirms that, if \mathcal{V}_c only contains m local minima, then it has m disconnected components. Then, if we focus on one of the minima, we can at the very least improve the size of its region of attraction’s estimate by “climbing” the energy’s level sets up to the first unstable equilibrium, and take the connected component to the minimum of \mathcal{V}_c as estimate.

While this improves the estimate, it still doesn’t completely clarify how to deal with more complicated situations. Consider for instance the energy function shown in Figure 1.6. The first sublevel set which contains an unstable equilibrium is \mathcal{V}_b , where b is the energy level of the saddle \mathbf{x}_b . While \mathcal{V}_b represents (the closure of) the actual controller-independent region of attraction of the global minimum, it doesn’t tell us anything about the two other minima. Moreover, even if we could remove \mathbf{x}_b and \mathbf{x}_c , the first unstable equilibrium which we encounter is now \mathbf{x}_f : the interior of the sublevel set \mathcal{V}_f has in-




Figure 1.6 – An example of energy function with non-trivial regions of attraction. Highlighted points represent its critical points, where green points are minima, yellow ones are saddles and the red one is a maximum. The letters are ordered according to the energy level.

deed three disconnected components, each one representing an (approximated) region of attraction for each minimum, but the estimate for \mathbf{x}_e is clearly pessimistic.

1.3.3 The connectivity tree: definition and examples

By restricting a bit the class of energy functions, we can actually go much further in this characterization. We propose a way to construct a graph, which we call the *connectivity tree*, that allows us to visualize how the regions of attraction of minima are interconnected. In particular, we want to show how we can use this tree as a graphical tool for deciding the ideal energy level to find the (component of the) sublevel set that best represents the region of attraction for a given minimum. We will also see that the connectivity tree will allow us to discuss a weaker form of the controller-independent region of attraction, which only guarantees us convergence with probability 1 to the minimum but could improve its size.

As a class of energy functions, we will consider functions at least $\mathcal{C}^2(\mathcal{M})$ with all nondegenerate equilibria (as always so far) and such that all sublevel sets are compact. Also, for simplicity, we will consider that the function has only a finite number of equilibria and that no two different equilibria have the same energy level. We will assume that all equilibria are known and that exact simulations can be made for an energy-decreasing system, at least for trajectories whose orbits are compact.

For didactic purposes, we will first show how to construct and use the connectivity tree for the example shown in Figure 1.6. We will make this construction more rigorous in the sequel, thanking the readers for this small act of faith.

The idea is simple: we want to construct a tree that represents the connectivity of the sublevel sets as the energy increases. Each leaf represents the creation of a new connected component, each node the junction of two components or the self-junction of a component, while each edge represents a connected component in the sublevel set, for all energy values between the ones of the two nodes connected by the edge. You can see a very simple example of a connectivity tree in Figure 1.7. Intuitively, we expect that the interior nodes represent the unstable critical points, and that the critical point with the highest energy is the root of the tree¹¹.

Let us now consider the function in Figure 1.6. We want to define an iterative algorithm that builds the connectivity tree from the ground up, relying on exact simulations to see

11. Assuming the manifold is connected; otherwise, we will have as many disconnected trees as there are components of the manifold.




Figure 1.7 – An example of connectivity tree for a simple energy function. Green points are minima while the yellow one is a saddle. As we climb the energy levels, the sublevel set firstly has a single component connected to the global minimum, then a second one connected to the local minimum, and finally they are re-connected once we reach the saddle’s energy level.

how the component of the sublevel sets are interconnected. We start with the lowest energy equilibrium (i.e., \mathbf{x}_a), which, being a minimum, creates a new component. We add our first leaf to the connectivity tree.

We then move up the energy lines, and the second equilibrium we encounter is \mathbf{x}_b , which is a 1-saddle. By simulating¹² its *heteroclinic orbits* (i.e., the two trajectories that form the saddle’s unstable manifold) we can see that they both converge to \mathbf{x}_a . Thus, at the energy level of \mathbf{x}_b , the sublevel set “folds” into itself, which doesn’t change its connectivity. We add a node to the connectivity tree, representing \mathbf{x}_b , and a single edge connecting it to \mathbf{x}_a .

Moving up, we get to \mathbf{x}_c , which is a 2-saddle. In this case, its unstable manifold is 2-dimensional, meaning that we cannot fully characterize it with simulations. However, we will see in the following that saddles of index > 1 cannot change the sublevel set’s connectivity, which in turn means that all trajectories in its unstable manifold converge to the same component of the sublevel set. Thus, we can identify this component by simulating a single trajectory, which, in our example, converges either to \mathbf{x}_a or \mathbf{x}_b . We

12. To simulate the heteroclinic orbits, we can use any energy-decreasing system. We can easily verify that they all lead to the same components by analyzing their unstable manifold in the chart defined by Morse lemma [74], restricted on a neighborhood of the saddle.




Figure 1.8 – An example of a nontrivial connectivity tree. Green points are minima, yellow ones are saddles and the red one is a maximum.

add the node corresponding to \mathbf{x}_c in our tree together with an edge connecting it to \mathbf{x}_b .

We can iterate this algorithm, inspecting all the remaining equilibria in order of energy level. Both \mathbf{x}_d and \mathbf{x}_e are minima, which correspond to isolated leaves on the tree. Then, by simulating the heteroclinic orbits of \mathbf{x}_f (which is a 1-saddle), we see that it links the components connected to \mathbf{x}_a and \mathbf{x}_d , and so the corresponding node has two edges connecting it to \mathbf{x}_c and \mathbf{x}_d . Analogously, the node corresponding to \mathbf{x}_g links \mathbf{x}_e and \mathbf{x}_f , completing the tree, which you can see in Figure 1.8.

We summarize the algorithm in Algorithm 1. The connectivity tree that we propose can be built in general for functions defined on n -dimensional manifolds, and, in our opinion, it is a useful graphical tool both for deducing the minima's controller-independent regions of attraction and for visualizing some of the function's properties that would otherwise be impossible to see. From the tree alone, we can immediately see that the best possible energy level to estimate the controller-independent region of attraction of \mathbf{x}_a is $b = v(\mathbf{x}_b)$; analogously, we can see that cutting the energy function just below f and g gives us estimates for \mathbf{x}_d and \mathbf{x}_e , respectively.

With the connectivity tree we can also discuss a weaker form of the region of attraction, which only guarantees convergence to the minimum *with probability 1*, i.e., we allow for a zero-measure subset of initial positions so that the corresponding trajectories do not converge to the minimum. For instance, if we take any sublevel set \mathcal{V}_c with $v(\mathbf{x}_b) < c < v(\mathbf{x}_f)$

Data: $\mathbf{x}_1, \dots, \mathbf{x}_m$, critical points of $v(\mathbf{x})$ s.t. $v(\mathbf{x}_1) < \dots < v(\mathbf{x}_m)$.
Result: (E, V) , edges and vertices of the connectivity tree.
 $E, V \leftarrow \emptyset, \{\mathbf{x}_1\}$
for $k = 2$ **to** n **do**
 $V \leftarrow V \cup \{\mathbf{x}_k\}$
 if $\text{index}(\mathbf{x}_k) == 1$ **then**
 /* Get highest equilibria within the components where the
 heteroclinic orbits converge to */
 $\mathbf{x}_a, \mathbf{x}_b \leftarrow \text{heteroclinic}(\mathbf{x}_k)$
 if $\mathbf{x}_a \neq \mathbf{x}_b$ **then**
 $E \leftarrow E \cup \{(\mathbf{x}_a, \mathbf{x}_k), (\mathbf{x}_b, \mathbf{x}_k)\}$
 else
 $E \leftarrow E \cup \{(\mathbf{x}_a, \mathbf{x}_k)\}$
 end
 else if $\text{index}(\mathbf{x}_k) > 1$ **then**
 /* Get highest equilibrium within the component where unstable
 orbits converge to */
 $\mathbf{x}_a \leftarrow \text{unstable}(\mathbf{x}_k)$
 $E \leftarrow E \cup \{(\mathbf{x}_a, \mathbf{x}_k)\}$
end

Algorithm 1: Pseudo-code for building the connectivity tree.

in the tree depicted in Figure 1.8, the component connected to \mathbf{x}_a does not approximate the region of attraction of \mathbf{x}_a , but still guarantees convergence with probability 1 to it, since the stable manifolds of \mathbf{x}_b and \mathbf{x}_c have measure zero. We believe that this weak region of attraction can be of great practical interest, since on the one hand, we could sensibly improve the region's size (as the considered example shows), while on the other, any amount of noise would prevent the system to converge to the unstable equilibria anyway.

1.3.4 The connectivity tree: formal justification

The inspiration for this construction comes from Morse theory, from which we will take in full force to justify the algorithm. There are two crucial properties underlying Algorithm 1, which are:

- Meeting a minimum creates a new connected component, and
- The connectivity of the sublevel set can change only when we encounter a saddle of index 1.

These two points justify that all leaves and nodes are critical points, and tell us that all unstable trajectories of a saddle of index > 1 converge to the same component, which motivates our choice to do only one simulation for these equilibria.

The main argument to justify these two properties can be derived from the following theorem, a classical result of Morse theory, which shows that the domain of a function such as those we consider has the same topological properties as a *CW-complex* constructed from the critical points of v . Intuitively, a CW-complex is a set \mathcal{E} of closed balls $e^k \subset \mathbb{R}^k$, $0 \leq k \leq n$, called *cells* or *k-cells*, together with a family of functions, called *characteristic maps*, that tell us how the balls “fit together”. Specifically, we require that the boundary of each cell must be connected in a “reasonable” way to lower-dimensional cells. For a formal introduction to Morse theory and CW-complexes, see [74] and [65], respectively.

Theorem 1.5 (Theorem 3.5 in [74]). *If v is a differentiable function on a manifold \mathcal{M} with no degenerate critical points, and if each of its sublevel sets is compact, then \mathcal{M} has the homotopy type of a CW-complex, with one cell of dimension k for each critical point of index k .*

We want to apply this theorem to each sublevel set of the energy function in order to discuss their connectedness. However, if we restrict the domain to an (open) sublevel

set of the energy function, then some sublevel sets of the restricted function will be non-compact in the subspace topology. Additionally, in order to be able to appropriately apply these results to the image-based visual servoing energy function, we want to generalize our analysis to functions whose sublevel sets are compact up to a certain energy level. Thus, we propose the following extension of Theorem 1.5 that allows for these relaxations. Intuitively, this result formalizes a simple geometrical intuition: we modify the function “stretching” it to infinity as we approach the boundary of the sublevel set, making all sublevel sets of the restricted function compact while retaining the energy function’s equilibria.

Corollary 1.2. *Let v be a differentiable function on a manifold \mathcal{M} and let $c \in \mathbb{R}$ be such that, for all $a < c$, the sublevel sets \mathcal{V}_a are compact and the critical points with energy level below c are nondegenerate. Then, the open sublevel set:*

$$\mathcal{M}^c = \{\mathbf{x} \in \mathcal{M} : v(\mathbf{x}) < c\} \quad (1.34)$$

has the homotopy type of a CW-complex, with one cell of dimension k for each critical point of index k within \mathcal{M}^c .

Proof. Let $\epsilon > 0$ be such that the set $v^{-1}([c - \epsilon, c])$ contains no critical points of v , and let $\rho : [0, c) \rightarrow [1, +\infty)$ be defined as:

$$\rho(x) = \begin{cases} 1, & 0 \leq x \leq c - \epsilon \\ 1 - (x - c + \epsilon) \log \left(1 - \frac{x - c + \epsilon}{\epsilon} \right), & c - \epsilon \leq x < c \end{cases}. \quad (1.35)$$

We can see the graph of this function in Figure 1.9, and it is easy to check that it has a vertical asymptote at $x = c$ and that it is smooth on its domain.

Since, up to c , the sublevel sets are compact, v has a global minimum by its continuity. We assume, without loss of generality, that this minimum is 0, so that the function $\bar{v}(\mathbf{x}) = \rho(v(\mathbf{x}))v(\mathbf{x})$ is well defined on \mathcal{M}^c . We have that all its sublevel sets are compact, since multiplying by $\rho(v(\mathbf{x}))$ does not change the shape of the sublevel sets, meaning that all sublevel sets of \bar{v} are equal to a sublevel set of v strictly below the energy level c , which is compact by hypothesis. Additionally, \bar{v} and v coincide on $\mathcal{V}_{c-\epsilon}$ (meaning that the critical points are the same), while, on $\mathcal{M}^c \setminus \mathcal{V}_{c-\epsilon}$, we have that:

$$\nabla \bar{v}(\mathbf{x}) = \underbrace{(\rho(v(\mathbf{x})) + v(\mathbf{x})\rho'(v(\mathbf{x})))}_{>0} \nabla v(\mathbf{x}), \quad (1.36)$$




Figure 1.9 – Graph of $\rho(x)$ with $c = 1$ and $\epsilon = 0.25$. The red dot denotes the point $(c - \epsilon, 1)$.

which implies that \bar{v} does not have any critical points in the region. Finally, \bar{v} is differentiable on \mathcal{M}^c (by the smoothness of ρ and the differentiability of v), meaning that we can apply Theorem 1.5, concluding the proof. \square

One of the critical information we can infer from the CW-complex is whether it is connected or not: a classical theorem of topology (see, e.g., [65, Proposition 5.11]) tells us that a CW-complex is connected if and only if its 1-skeleton (i.e., the subcomplex formed only by 0- and 1-cells) is connected. This characterization justifies our choice to focus only on minima and 1-cells in drawing the connectivity tree. As we move up the energy values, when we encounter a minimum we add a 0-cell to the sublevel set (i.e., we add an isolated point, representing a new component), while when we encounter a 1-saddle we add a 1-cell (which we visualize as an arc) that either joins two components or self-joins the same component. These two elements form the “building blocks” for the connectivity tree, to which we finally add the (> 1)-index saddles (which do not impact connectivity) that allow us to distinguish between regions of attraction in the strict sense or with probability 1.

1.3.5 Detection and discrimination of regions of attraction

We end the section with a brief discussion on a practical issue in using our strategy with high-dimensional systems: given an element in a sublevel set with several components, how do we distinguish to which one the element belongs to? If we are able to provide a way to discriminate the components, we could, on the one hand, extract more information from them (e.g., compare their volumes), while, on the other, we could equip the system

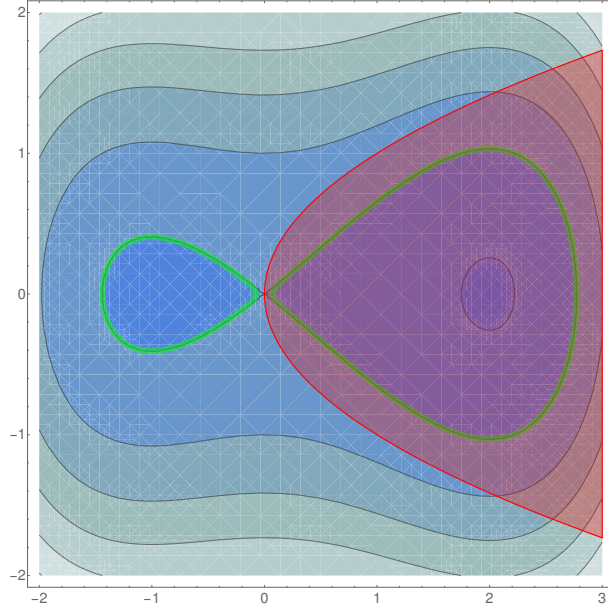


Figure 1.10 – An example of components’ separation in \mathbb{R}^2 . We show the level sets of $v(\mathbf{x}) = 0.1x_1^2(x_1^2 - \frac{4}{3}x_1 - 4) + x_2^2$, and we separate the sublevel set $\mathcal{V}_{-\epsilon}$, whose boundary is depicted in green. The separator function is $g(\mathbf{x}) = x_1 - x_2^2$, and we depict the region defined by $g(\mathbf{x}) \geq 0$ in red.

with a way to detect whether it fell into a controller-independent region of attraction. This last possibility would also be, we believe, a first step in proposing control strategies that allow to escape an undesired stable equilibrium.

Given a sublevel set \mathcal{V}_c with m components, each of which we will denote with $\mathcal{V}^{(i)}$, a possible approach to identify a specific component $\mathcal{V}^{(0)}$ is to determine $(m - 1)$ functions $g_i : \mathcal{M} \rightarrow \mathbb{R}$ such that each g_i is strictly negative on $\mathcal{V}^{(0)}$ and strictly positive on $\mathcal{V}^{(i)}$; said differently, the functions $g_i(\mathbf{x})$ are such that each level set $\{\mathbf{x} \in \mathcal{M} : g_i(\mathbf{x}) = 0\}$ separates $\mathcal{V}^{(0)}$ and $\mathcal{V}^{(i)}$. You can see a very simple example of a separation of 2 components in Figure 1.10. With this family of functions, we can easily compute whether a point \mathbf{x} with $v(\mathbf{x}) \leq c$ is within $\mathcal{V}^{(0)}$ by checking the signs of $g_i(\mathbf{x})$, and we can also characterize $\mathcal{V}^{(0)}$ as the intersection of \mathcal{V}_c and all the sublevel sets $\mathcal{G}_0^{(i)} = \{\mathbf{x} \in \mathcal{M} : g_i(\mathbf{x}) \leq 0\}$.

Unfortunately, there is no obvious way to identify these functions, especially when the state-space is not Euclidean, and, at the current time, we only have partial answers on how to deal with this problem. We present a simple strategy that approximates the separating functions g_i with hyperplanes, which we can always define under the assumption that the state-space is \mathbb{R}^n . Since we cannot expect that any two components can be separated by

hyperplanes, we also allow for lowering the energy level: in this way, the components of the sublevel set will shrink, resulting in more “leeway” in the choice of the hyperplane. This always works if we want to identify the component connected to the global minimum, since, in the worst case scenario, the other components eventually disappear as we decrease the energy. For this reason, however, we cannot guarantee that this strategy lets us always identify other components in general, even though we expect it to work satisfactorily on average cases.


We propose the following heuristic to identify the separating functions. Let $c \in \mathbb{R}$ be the energy level we are interested in, assume that \mathcal{V}_c has m components and let $\mathbf{m}_0, \dots, \mathbf{m}_{m-1}$ be the minimizers within each component. Since the function g_i must, at the very least, allow us to discriminate between the minimizers within $\mathcal{V}^{(0)}$ and $\mathcal{V}^{(i)}$, we consider as a candidate the hyperplane which is perpendicular to the segment connecting \mathbf{m}_0 and \mathbf{m}_i and which passes through the point that maximizes the energy on this segment, which we denote with \mathbf{p}_i . Then, in order to guarantee the separation of the components, we evaluate the minimal energy $c_i \leq c$ on the hyperplane and take this energy level as the new c . In formulae, we solve these two optimization problems:

$$\begin{cases} \mathbf{p}_i = \arg \max_{\mathbf{x}=t\mathbf{m}_0+(1-t)\mathbf{m}_i, t \in [0,1]} v(\mathbf{x}), \\ c_i = \min_{g_i(\mathbf{x})=0} v(\mathbf{x}) \end{cases} \quad (1.37)$$


Then, we iterate over all the components to find all the separating functions and finally identify the new energy level $\bar{c} = \min\{c_1, \dots, c_{m-1}\}$ that allows us to isolate $\mathcal{V}^{(0)}$.

We can see an example of such a splitting in Figure 1.11, where we considered a sublevel set immediately below the energy level of the 1-saddle \mathbf{x}_d , for which we wanted to identify the component connected to the global minimum. We can clearly see in Figure 1.11a that this sublevel set has three components, one for each minimum. Using the strategy described above, we were able to identify two straight lines that separate the three components, with only a minimal decrease in energy level. Thanks to these separating function, we were able to explicitly compute $\mathcal{V}^{(0)}$, which is shown in green in Figure 1.11b.

It is easy to see some steps that we can take to improve the estimates provided by this procedure. For instance, we could frame the two optimization problems (1.37) as a single *maximin* problem, in order to find the point \mathbf{p}_i that maximizes the energy’s minimum on the hyperplane. Other improvements can be obtained by parameterizing the hyperplanes,



(a) Energy function.



(b) Separation of the sublevel set's components.

Figure 1.11 – Separation of the sublevel set's components with the procedure described in Sec. 1.3.5. The local minima are depicted in green, while the saddle are in yellow. The component connected to the global minimum is computed as $\mathcal{V}^{(0)} = \mathcal{V}_c \cap \mathcal{G}_0^{(1)} \cap \mathcal{G}_0^{(2)}$.

or selecting higher-order polynomials as candidates, and including these parameters in the optimization process. Nevertheless, we decided to settle for the heuristic proposed above because we expect it to have a computational cost which is comparable to the one of computing all critical points of the energy function, since, as we will see, this computation alone can be barely tractable in practice.

1.4 Non-compact sublevel sets in the visual servoing context

Throughout this chapter, we investigated several interesting properties of what we call energy-decreasing systems, with the final objective of applying these ideas to image-based visual servoing systems. Our arguments above on using the energy function's sublevel sets as estimates of the regions of attraction rely on the application of LaSalle invariance principle to show that a given sublevel set of the energy function is (positively) invariant for any energy-decreasing system, and, furthermore, the construction of the connectivity tree relies heavily on the Morse theorem (Theorem 1.5). However, both these important results require that the sublevel sets are compact, which, as we will see, it is not always the case for the energy function that we consider in image-based visual servoing.

It is possible to identify two significant behaviors, which showcase that some sublevel sets are not compact. The first one is well known: we can find examples where, for some trajectories, the camera's position moves back towards infinity, even though an energy-decreasing controller is employed. This means that the energy level can be bounded as we approach infinity, implying that some level sets are unbounded, and, thus, non-compact. The second behavior is more subtle, and, apparently, unknown in the literature: it is possible for the camera's position to approach one of the 3D points (i.e., make one of the depths Z_i converge to zero) while keeping a bounded energy level, and, in fact, we were able to identify some edge cases, in the context of planar visual servoing, where energy-decreasing trajectories converge towards one of the points. If we consider the state-space of the camera to be the set of $\mathbf{x} \in \text{SE}(3)$ such that all the points are in the camera's field of view, we will see that this behavior leads to the surprising conclusion that some sublevel sets are non-compact because of the state-space's topology.

The strategy that we propose to manage the non-compactness is to characterize the minimal energy level at which it arises and restrict our analysis to the sublevel sets which are below this critical level. By taking this set as the camera's state-space, we can




Figure 1.12 – Example of critical energy level represented in the connectivity tree. The function has an horizontal asymptote to the right, meaning that all sublevel sets above the red square value will be non-compact.

safely discuss the convergence properties of any energy-decreasing system, at the cost of “giving up” discussing what happens outside. We will see that, for a representative class of examples in visual servoing, we can always find such a sublevel set which contains all the critical points of interest, indicating that these problems do not prevent, in general, to carry out a meaningful analysis of these systems.

In order to highlight the critical energy levels at which the sublevel sets become non-compact, we include them as well in the connectivity tree. We can see an example of this in Figure 1.12, where we consider a simple scalar function which presents an horizontal asymptote as x tends to infinity. The critical energy level, represented by a red square in the figure, marks the “safety limit” below which we can still discuss the convergence properties. We will see this construction again in Chapter 3, where we use it to represent the critical level sets discussed in this section for a set of image-based visual servoing examples.

In the following, we will study the two situations separately. For both of them, we first provide simple examples that let us visualize the issues in the comforting settings of \mathbb{R}^2 , and we then derive a strategy to compute their critical energy level.

1.4.1 Non-compactness of sublevel sets due to unboundedness

The first problem we tackle is to find the minimal energy level as the camera runs off to infinity. Energy-decreasing systems in Euclidean settings with this behavior are very well understood in the Lyapunov theory literature: if the Lyapunov function is not radially




Figure 1.13 – Example of a Lyapunov function which is not radially unbounded.

unbounded, then its sublevel sets are unbounded and error decreasing trajectories may escape to infinity. A canonical example is given in [60], where we consider the Lyapunov function:

$$v(x) = \frac{x_1^2}{1+x_1^2} + x_2^2. \quad (1.38)$$

As we can also see in Figure 1.13, $v(x)$ is not radially unbounded in the direction $u = (1, 0)$, since $v(tu) = \frac{t^2}{1+t^2}$ is bounded. Here we expect the gradient descent $\dot{x}(t) = -\nabla v(x(t))$ to be globally asymptotically stable, but there exist some energy decreasing trajectory that escape to infinity. Consider, for instance, $x(t) = (t, \frac{1}{t})$: the time derivative of $v(x)$ along this trajectory is $\frac{d}{dt}v(x(t)) = -\frac{2(1+2t^2)}{t^3(1+t^2)^2}$, which is strictly negative for $t > 0$.

As discussed above, a well-known example in image-based visual servoing is representative of this situation¹³: observing the Cartesian coordinates of 4 points, an escaping trajectory is obtained by arranging the points in a square configuration and choosing a desired pose such that the visual axis is perpendicular to the square and centered with respect to it, and pick a starting pose with 180° axial rotation with respect to the desired one [20]. Furthermore, we will show (in Chapter 2) that this system has one unique fixed point, which is the global minimum of the energy function, which illustrates that such

13. It's important to notice that, in general, boundedness is not properly defined on manifolds. However, by seeing $SE(3) = \mathbb{R}^3 \times SO(3)$, we can indirectly use it to tell whether a subset $S \in SE(3)$ is compact: if its natural projection onto \mathbb{R}^3 is unbounded (i.e., not compact), then S cannot be compact, since the natural projection is continuous [65, Corollary 3.28] and the image of a compact set under a continuous function is compact [65, Theorem 4.32].

escaping trajectories are not detectable by looking only at fixed points.

We now show how we can formally derive the minimal energy level at infinity. We begin by noticing that, as the camera moves away from the tracking points, the features $\mathbf{s}_i = (x_i, y_i)$ tend to get closer to each other. By the geometry of the problem, it's easy to convince ourselves that each point's features all converge to an identical value (potentially infinite) when the camera is infinitely far from the tracking points. We denote with $\bar{\mathbf{s}}$ this limit value, i.e.:

$$\lim_{\|\mathbf{o}\mathbf{t}_c\| \rightarrow \infty} \mathbf{s}_i = \bar{\mathbf{s}}, \quad i = 1, \dots, N. \quad (1.39)$$

It's then possible to rewrite the energy function (1.5) as a function of $\bar{\mathbf{s}}$, as:

$$v(\bar{\mathbf{s}}) = \frac{1}{2} \sum_{i=1}^N \|\bar{\mathbf{s}} - \mathbf{s}_i^*\|^2, \quad (1.40)$$

where \mathbf{s}_i^* denotes the desired features for point i . We can find the minimum of $v(\bar{\mathbf{s}})$ by checking where its gradient, taken with respect to $\bar{\mathbf{s}}$, is null. We have that:

$$\nabla v(\bar{\mathbf{s}}) = \sum_{i=1}^N \bar{\mathbf{s}} - \mathbf{s}_i^* = 0 \iff \bar{\mathbf{s}} = \frac{1}{N} \sum_{i=1}^N \mathbf{s}_i^* \quad (1.41)$$

This critical value of $\bar{\mathbf{s}}$ must be a minimum, being $v(\bar{\mathbf{s}})$ a sum of strictly convex functions and, then, convex itself. Thus, we identified the minimal energy level c_∞ such that the sublevel set are unbounded, which is, simply:

$$c_\infty = \frac{1}{2} \sum_{i=1}^N \left\| \left(\frac{1}{N} \sum_{i=1}^N \mathbf{s}_i^* \right) - \mathbf{s}_i^* \right\|^2 \quad (1.42)$$

1.4.2 Non-compactness of sublevel sets due to the state-space's topology

As we already hinted here and there, compactness on manifolds does not always agree with the intuition we have from Euclidean space. For instance, we might have a sublevel set which is closed and bounded (for a given metric on the manifold), but non-compact. We provide here an example that lets us introduce the second problematic behavior we encounter in the context of image-based visual servoing, i.e., that the error may remain finite if one point converges to the camera center while its depth tend to zero. Points' features in the screen are foreseen to diverge to infinity when their depth tends to zero,




Figure 1.14 – Graph (left) and stream plot (right) of an attractive (top) and a repulsive (bottom) hole-in-the-wall.

and, thus, the error is foreseen to diverge to infinity as well, enforcing compactness of the sublevel set. But, as a consequence of the energy's boundedness when the camera position is superposed to a point, this wall of infinite error at zero depth contains a "hole" where the error remain finite. We call this property a *hole-in-the-wall*. The following two energy functions $v_+(\mathbf{x})$ and $v_-(\mathbf{x})$ represent two-dimensional instances of holes-in-the-wall:

$$v_{\pm}(\mathbf{x}) = \frac{x_2^2}{x_1^2} + (x_1 \pm 2)^2 + x_2^2, \quad (1.43)$$

where we consider the state space to be $\mathcal{X}_+ = \{\mathbf{x} \in \mathbb{R}^2 : x_1 > 0\}$. In both cases, when $x_2 \neq 0$ the energy function presents an infinite wall at $x_1 = 0$, hence any error decreasing controller will tend to increase small values of x_1 . However, $v_{\pm}(x_1, 0) = (x_1 \pm 2)^2$, hence it presents a finite limit value when x_1 tends to 0 with $x_2 = 0$. The graph and the phase portrait of the gradient flow are shown in Figure 1.14. In the case of v_+ , the energy in the hole-in-the-wall is the lowest, and all trajectories converge toward it. Since the vector field is not defined there, the trajectories actually stop in finite time. Noticeably, all nonempty sublevel sets are bounded for the metric inherited from \mathbb{R}^2 ; additionally, being v_+ continuous on \mathcal{M}_+ , the preimage of a closed set is closed in the subspace topology, meaning that the sublevel sets are closed as well. However, they are clearly not positively invariant for the gradient flow: thus, by LaSalle invariance principle, they cannot be compact.

In the case of v_- , the energy in the hole-in-the-wall is not a local minimum, and all trajectories seem to converge towards the global minimum $(2, 0)$. In this case, we cannot exploit LaSalle invariance principle. We can use a different, more direct characterization of compactness on manifolds: we have that any *sequentially compact* set is compact [65, Theorem 4.45], i.e., a set is compact if and only if every sequence of points in the set has a subsequence that converges to a point in the set. Thus, it is clear that sublevel sets that are above the energy of the hole-in-the-wall are not compact, since we can always pick a sequence of points in the sublevel set that converges to the origin in \mathbb{R}^2 , but does not have any converging subsequence in \mathcal{X}_+ . On the other hand, all lower sublevel sets are compact.

Remark. *Intuitively, we could think about compactness on manifolds in the following terms: a closed¹⁴ subset of a manifold \mathcal{M} is not compact if there is a sequence $\{\mathbf{x}_k\}_{k=1}^{+\infty} \in$*

14. For general topological spaces, a compact set doesn't need to be closed: for instance, any proper subset of a finite, discrete topological space endowed with the trivial topology is compact (since any

\mathcal{M} that tries to “run off” to a point outside the manifold, in the same way of a closed, unbounded set which is not compact because there is a sequence that runs off to infinity. As an example, we can compare \mathcal{X}_+ (as defined above) and \mathbb{R}^2 , which are homeomorphic under the homeomorphism $f(\mathbf{x}) = [\log x_1, x_2]$. Since it is continuous, $f(\mathbf{x})$ maps compact sets to compact sets. However, any subset of \mathcal{X}_+ which touches its boundary, even if closed and bounded, gets mapped to a unbounded set of \mathbb{R}^2 .

These holes-in-the-wall have been unnoticed in the visual servoing literature until now, most probably because their seemingly little practical interest. The dynamics in their vicinity is not understood, and, in fact, some examples of image-based visual servoing systems can be constructed in such a way that trajectories converge toward a hole-in-the-wall. As discussed above, as a first step towards the investigation of these systems’ global stability, we want to characterize their energy level in order to investigate the behavior of the system below this energy level.

Contrarily to the previous section, we will not be able to derive a closed-form solution that represents the minimal energy at the holes-in-the-wall, but we will instead frame it as an optimization problem, and explicitly write the first-order optimality conditions. Assume that the camera is superposed to the i -th point ${}^o\mathbf{a}_i$ and that all the other points are in front of the camera (which is always possible as long as the i -th point is not in the convex hull of the others). The energy function (1.5) has a jump discontinuity at this position, since the visual error \mathbf{e}_i corresponding to ${}^o\mathbf{a}_i$ is not well defined but its limit as the camera position approaches ${}^o\mathbf{a}_i$ is, for any orientation such that all the points are kept in front of the camera. Since this depends on the camera position alone, we actually have a 3-dimensional continuum of discontinuities, due to the 3 degrees of freedom given by the choice of orientation.

As a visual aid, we can see in Figure 1.15 a lower dimensional representation of this situation, where we clumped together the position coordinates on axis \mathbf{t} and the rotation ones on axis \mathbf{R} . Our objective is to identify the orientation that minimizes the energy along the discontinuity (represented in red in the Figure). The strategy we propose is to consider only the lower part of the jump and minimize it as a function of \mathbf{R} alone. Since, as we said, all the errors except \mathbf{e}_i are well defined on the discontinuity, we can trivially

open cover of the subset must be finite) but not closed. However, we assume that manifolds are always Hausdorff by definition, which implies that closedness is necessary for compactness [65, Proposition 4.36].




Figure 1.15 – Simplified representation of a hole-in-the-wall in image-based visual servoing. When the camera position \mathbf{t} is fixed on ${}^o\mathbf{a}_i$, the energy function is discontinuous for any orientation. The choice of orientation that minimizes the energy on the discontinuity is shown in red.

rewrite the energy function as:

$$v(\mathbf{x}) = \sum_{k \neq i}^N \|\mathbf{e}_k(\mathbf{x})\|^2 + \|\mathbf{e}_i(\mathbf{x})\|^2. \quad (1.44)$$

Thus, it is clear that the lower part of the jump is obtained by putting $\mathbf{e}_i(\mathbf{x}) = 0$, if we assume that, for any given orientation \mathbf{R} , we can always find a sequence of camera poses that converges to $({}^o\mathbf{a}_i, \mathbf{R})$ as \mathbf{e}_i converges to zero (we'll verify this assumption shortly). We finally have all the ingredients we need to find the minimal energy level at a hole-in-the-wall: assuming that $\mathbf{e}_i = 0$ and fixing the position of the camera at ${}^o\mathbf{a}_i$, i.e., taking $\mathbf{x} = ({}^o\mathbf{a}_i, \mathbf{R})$, the energy function (1.44) becomes a smooth function of \mathbf{R} , meaning that we can take its gradient and find the minimum as one of the solutions of $\nabla v(\mathbf{R}) = 0$. Finally, since each point is a (possible) hole-in-the-wall, we perform the computation for each of them and just keep the minimum among the energy levels.

As we said above, once we find the orientation that minimizes the energy, we need to find a sequence of poses that converge to ${}^o\mathbf{a}_i$ as \mathbf{e}_i goes to zero, which allows us to verify that we actually obtained a minimum and not a lower bound for it. Actually, we

can easily build a sequence of camera poses that converges to ${}^o\mathbf{a}_i$ while keeping the error $\mathbf{e}_i(\mathbf{x})$ null throughout the whole sequence. Let the camera be superposed to ${}^o\mathbf{a}_i$, meaning that ${}^c\mathbf{a}_i = \mathbf{0}$, and consider the vector $\mathbf{n} = (x_i^*, y_i^*, 1)$ expressed in camera frame. If we move the camera in the direction $-\mathbf{n}$ while keeping the orientation constant, then the point's features will be $(x_i, y_i) = (x_i^*, y_i^*)$ for any non-null position displacement, implying $\mathbf{e}_i(\mathbf{x}) = 0$. Additionally, we can easily see that, as the camera moves, all the points (expressed in camera frame) undergo a translation in the direction of \mathbf{n} , meaning that they are all kept in front of the camera (since, by hypothesis, all the points have positive depths at the minimal orientation). Thus, any sequence of poses such that the positions are taken along this direction and the orientation is fixed to the minimal one converges to ${}^o\mathbf{a}_i$ keeping the error $\mathbf{e}_i(\mathbf{x})$ null, regardless of the minimizing orientation.

In practice, this optimality condition is very complicated to solve, but we will see in the next chapter a change of variables that renders the system of equations tractable. Since the space of orientations is compact, we will be able to use an interval analysis-based solver which allows us to find all the critical points of $v(\mathbf{R})$ in a verified way, guaranteeing that the global minimum is found.

COMPUTATION OF THE IMAGE-BASED VISUAL SERVOING ENERGY'S CRITICAL POINTS

In this chapter, we will apply all the strategies that we developed in Chapter 1 to a class of energy-decreasing image-based visual servoing systems. As we will see, the main difficulty in studying these systems is represented by the verified computation of all the critical points of their energy function (i.e., their equilibria), which will take up most of the chapter.

The chapter is organized as follows. In Section 2.1 we consider again three well-known energy-decreasing image-based visual servoing controllers that use the Cartesian coordinates in the image of N points as visual features, namely the pseudo-inverse, transpose and Levenberg-Marquardt controllers. For these systems, we investigate the natural parametrization of the equilibrium condition and we argue that, at the current state, calculating all its solutions is out of reach of current complete solvers for systems of equations. Then, we propose a different strategy to compute the equilibria directly in features space. Solving the equilibrium condition as a function of the features opens up questions regarding the reconstruction of the camera pose and its uniqueness for a given set of features, questions which are addressed in Section 2.2. In Section 2.3, we present a strategy that we can use to rigorously find all the energy function's critical points in practice, namely a computational algebra-based method that rely on *Gröbner bases* to efficiently solve systems of polynomial equations.

The equilibrium condition developed in this chapter can be used for other classes of image-based visual servoing controllers, even ones which are not energy decreasing. We will discuss one such class of controllers in Chapter 4, where we will show that their equilibrium condition can then be seen as a (significantly simpler) particular case of the energy-decreasing one.

Table 2.1 – Equilibrium condition's explicit expression.

$$\sum_{i=1}^N \left[\left(\prod_{j \neq i}^N Z_j \right) (x_i^* - x_i) \right] = 0 \quad (2.2a)$$

$$\sum_{i=1}^N \left[\left(\prod_{j \neq i}^N Z_j \right) (y_i^* - y_i) \right] = 0 \quad (2.2b)$$

$$\sum_{i=1}^N \left[\left(\prod_{j \neq i}^N Z_j \right) \left(x_i(x_i - x_i^*) + y_i(y_i - y_i^*) \right) \right] = 0 \quad (2.2c)$$

$$\sum_{i=1}^N x_i y_i (x_i - x_i^*) + (1 + y_i^2)(y_i - y_i^*) = 0 \quad (2.2d)$$

$$\sum_{i=1}^N (1 + x_i^2)(x_i^* - x_i) + x_i y_i (y_i^* - y_i) = 0 \quad (2.2e)$$

$$\sum_{i=1}^N (x_i y_i^* - y_i x_i^*) = 0. \quad (2.2f)$$

2.1 Mathematical models for the visual servoing equilibrium condition

As we saw in the previous chapter, the gradient of the energy function that we have in Cartesian image-point-based visual servoing is $\mathbf{L}_e^\top [\mathbf{s}(\mathbf{x}) - \mathbf{s}^*]$, where the explicit expression of the interaction matrix \mathbf{L}_e is given in (1.2). Thus, the strategy that we propose to compute all the critical points is simply to solve:

$$\mathbf{L}_e^\top [\mathbf{s}(\mathbf{x}) - \mathbf{s}^*] = \mathbf{0}. \quad (2.1)$$

However, as we will see in Section 2.3, in order to be able to solve (2.1) in practice we need to find a *polynomial* system of equation which has the same solution set. Given the interaction matrix structure, it is easy to *polynomialize* (2.1), whose explicit expression is given in (2.2) (Table 2.1). In (2.2), we simply multiplied the first three equations by the product of all Z_i 's to make the expression polynomial. Since the system's state-space is such that all the points depths Z_i are strictly positive, this transformation does not change the solution set. However, (2.2) is expressed in what is called the *extended features* $\boldsymbol{\xi} =$

$(\mathbf{s}, \mathbf{Z}) \in \mathbb{R}^{3N}$ [92], meaning that it only provides a necessary condition for the equilibrium: it is a system of 6 equations in $3N > 6$ unknowns, where *infeasible extended features* (i.e., extended features not corresponding to any camera pose) are equally considered together with *feasible extended features*. Thus, it is imperative to find a strategy that allows us to impose that the solutions of (2.2) are within the set of feasible extended features. We propose two such strategies, which will be discussed in the following.

2.1.1 Camera state-space representation

A first attempt to solve the equilibrium condition is to consider the camera pose ${}^c\mathbf{T}_o \in \text{SE}(3)$ as the system's variables, from which features $\mathbf{s}({}^c\mathbf{T}_o)$ and their corresponding depths $\mathbf{Z}({}^c\mathbf{T}_o)$ are easily computed, leading to intrinsically feasible extended features that can be substituted in (2.2). However, $\text{SE}(3)$ does not admit a unique representation. We focus here on two of these representations, namely, rotation matrices and quaternions. We then rewrite (2.2), which is naturally expressed as a function of $\boldsymbol{\xi}({}^c\mathbf{T}_o)$, for the two parameterizations. First, considering the parameterization ${}^c\mathbf{T}_o = ({}^c\mathbf{R}_o, {}^c\mathbf{t}_o)$, we obtain the points coordinates in camera frame by

$${}^c\mathbf{a}_i = {}^c\mathbf{R}_o {}^o\mathbf{a}_i + {}^c\mathbf{t}_o, \quad (2.3)$$

from which the image features can be retrieved using (1.1) and their depth from the third component of ${}^c\mathbf{a}_i$. Second, this relation also allows us to rewrite (2.2) as a function of the quaternion parameterization ${}^c\mathbf{T}_o = ({}^c\mathbf{q}_o, {}^c\mathbf{t}_o)$, by expressing ${}^c\mathbf{R}_o$ as a function of ${}^c\mathbf{q}_o$ (see, e.g., [90]).

In order to actually solve the system on $\text{SE}(3)$, it is additionally necessary to impose the group constraints ${}^c\mathbf{R}_o \in \text{SO}(3)$ or ${}^c\mathbf{q}_o \in \mathbb{H}$ for the two representations. Putting all together, the explicit expression of (2.2) as a function of the camera state becomes a huge degree $3N$ polynomial system of 12 equations in 12 variables with the $({}^c\mathbf{R}_o, {}^c\mathbf{t}_o)$ -representation, while we have a degree $6N$ polynomial system of 7 equations in 7 variables with the $({}^c\mathbf{q}_o, {}^c\mathbf{t}_o)$ representation. A comparison between the degrees of these systems is detailed in Table 2.2. Because of their high degrees and substantial density, these systems are out of reach of current state-of-the-art complete solvers for polynomial systems. Note that similar conclusions are obtained using minimal representations of rotations, such as the angle multiplied by the unitary rotation axis.

Table 2.2 – Complexity comparison between different choices of parametrization for the perfect approximation case.

N	Parametrization	# of unknowns	# of equations	Degrees
N	$({}^c\mathbf{R}_o, {}^c\mathbf{t}_o)$	12	12	$[2N - 1, 2N - 1, 3N - 1, 3N, 3N, 2N - 1, 2, \dots, 2]$
	$({}^c\mathbf{q}_o, {}^c\mathbf{t}_o)$	7	7	$[2(2N - 1), 2(2N - 1), 2(3N - 1), 6N, 6N, 2(2N - 1), 2]$
	ξ	$3N$	$6 + \frac{N(N-1)}{2} + 1$	$[N, N, N + 1, 3, 3, 1, 4, \dots, 4, 5]$
$N = 4$	$({}^c\mathbf{R}_o, {}^c\mathbf{t}_o)$	12	12	$[7, 7, 11, 12, 12, 7, 2, \dots, 2]$
	$({}^c\mathbf{q}_o, {}^c\mathbf{t}_o)$	7	7	$[14, 14, 22, 24, 24, 14, 2]$
	ξ	12	13	$[4, 4, 5, 3, 3, 1, 4, \dots, 4, 5]$
$N = 5$	$({}^c\mathbf{R}_o, {}^c\mathbf{t}_o)$	12	12	$[9, 9, 14, 15, 15, 9, 2, \dots, 2]$
	$({}^c\mathbf{q}_o, {}^c\mathbf{t}_o)$	7	7	$[18, 18, 28, 30, 30, 18, 2]$
	ξ	15	17	$[5, 5, 6, 3, 3, 1, 4, \dots, 4, 5]$
$N = 6$	$({}^c\mathbf{R}_o, {}^c\mathbf{t}_o)$	12	12	$[11, 11, 17, 18, 18, 11, 2, \dots, 2]$
	$({}^c\mathbf{q}_o, {}^c\mathbf{t}_o)$	7	7	$[22, 22, 34, 36, 36, 22, 2]$
	ξ	18	22	$[6, 6, 7, 3, 3, 1, 4, \dots, 4, 5]$

2.1.2 Extended features representation

As explained just above, expressing the system directly in SE(3) leads to a highly complicated system, which cannot be solved as-is. The equilibria condition (2.2), considered directly in the extended features space, is seemingly far simpler than the system's expression in camera state-space. We aim here to replace the extended feature feasibility, expressed in the previous system using the camera pose, by simpler constraints acting only on the extended features.

To this end, we propose to check the feasibility of the extended features ξ by first reconstructing the corresponding points in the camera frame by using the one-to-one correspondence

$$Z_i = Z_i, X_i = x_i Z_i, Y_i = y_i Z_i, \quad (2.4)$$

and second imposing that point-to-point inter-distances in the camera frame to be the same as in the world frame. This leads to the following $C_2^N = \frac{1}{2}N(N-1)$ point-to-point inter-distance constraints:

$$(x_i Z_i - x_j Z_j)^2 + (y_i Z_i - y_j Z_j)^2 + (Z_i - Z_j)^2 = d_{ij}^2, \quad (2.5a)$$

for all $i < j = 1, \dots, N$, where the distances $d_{ij} := \|\mathbf{a}_i - \mathbf{a}_j\|$ are known *a priori*. The new system (2.2) and (2.5a), which allows us to compute the system's equilibria directly in extended features space, has much lower degrees with respect to the camera state-space representation (see Table 2.2) and it is in the scope of current complete solvers for polynomial systems, as shown in Section 2.3.

This new expression offers a great simplification, but also comes with a drawback: the point-to-point inter-distance constraints are only a necessary condition that is not sufficient for an extended feature to be feasible. Informally, points in the camera frame that have the same point-to-point inter-distance as in the world frame may come from an improper rigid transformation belonging to the group of isometries E(3) and not to SE(3).

It is possible to overcome this issue by noting that the handedness of a tetrahedron is only preserved under *proper* isometries. Therefore, the following constraint is satisfied

only by *feasible* extended features:

$$\det \begin{bmatrix} x_1 Z_1 & x_2 Z_2 & x_3 Z_3 & x_4 Z_4 \\ y_1 Z_1 & y_2 Z_2 & y_3 Z_3 & y_4 Z_4 \\ Z_1 & Z_2 & Z_3 & Z_4 \\ 1 & 1 & 1 & 1 \end{bmatrix} = \det \begin{bmatrix} {}^o\mathbf{a}_1 & \cdots & {}^o\mathbf{a}_4 \\ 1 & \cdots & 1 \end{bmatrix}, \quad (2.5b)$$

where $\{{}^o\mathbf{a}_i\}_{i=1}^4$ are any 4 points in the configuration and where the two determinants are proportional to the *signed volume*¹ of the tetrahedron represented in the camera-frame and world-frame, respectively.

In the non-coplanar case, by taking 4 non-coplanar points in (2.5b), the *feasibility constraints* (2.5) select only the feasible extended features. In the coplanar case, handedness loses meaning due to the tetrahedron’s degeneracy and the distance constraints (2.5a) become sufficient for feasibility. Nevertheless, the addition of (2.5b) still has the benefit of speeding up the computations: informally, extended features corresponding to proper and improper isometries are merged in the coplanar case, leading to solutions of multiplicity two, and the determinant constraint removes this multiplicity.

2.2 Pose reconstruction

As we saw in the previous section, formulating the equilibrium condition directly in the extended features space gives us a system of equations which is within the reach of state-of-the-art complete solvers for polynomial systems. However, the solutions we find are extended features, while, in practice, it is crucial to find the corresponding camera pose. In particular, it is not a priori evident if feasible extended features only have one corresponding camera pose, or if several camera poses can lead to the same extended features.

In this Section, we first provide a reconstruction algorithm based on a closed-form solution of the *orthogonal Procrustes’ problem* [50]. We then give an original result showing that, in general, the pose reconstruction might be non-unique, and explicitly derive the set of all solutions of Procrustes’ problem. As a corollary, we show that, in the context of image-based visual servoing, there is always a unique solution. Finally, we provide a few examples, not coming from visual servoing, where we unexpectedly have a continuum of

1. By its definition, the absolute value of the determinant is 6 times the volume, and the sign of the determinant corresponds to the handedness.

solutions.

2.2.1 Reconstruction algorithm

We discuss here a classical pose reconstruction strategy and we adapt it to recover the camera pose given the extended features ξ and the points' coordinates in world frame $\{^o\mathbf{a}_i\}_{i=1}^N$, depending on whether the points are coplanar or not. The proposed reconstruction algorithm is based on the solution of the orthogonal Procrustes problem, whose aim, given two sets of points $\{\mathbf{P}_i\}_{i=1}^N$ and $\{\mathbf{Q}_i\}_{i=1}^N$, is to find an orthogonal matrix Ω that maps \mathbf{P}_i to \mathbf{Q}_i for all i in the best possible way. The problem allows for the points to undergo a translation as well; in this case, the translation is first removed from the problem by shifting the points' centroid to the origin, and then reconstructed using Ω and the centroids' information.

The reconstruction algorithm for a given solution ξ of (2.2) and (2.5) is detailed in Table 2.3, where the different treatment due to the coplanarity is enclosed in Step 3. The feasibility of ξ implies that there always exists at least one proper camera pose that corresponds to it, or, equivalently, that there exists at least one proper isometry between the two sets of points \mathbf{P} and \mathbf{Q} . Such isometry is a global minimizer of the residual $\|\Omega\mathbf{P} - \mathbf{Q}\|_F$, meaning that it is retrieved by the reconstruction algorithm, being the reconstructed orientation guaranteed to minimize the residual (see [56, 57]). However, since the algorithm computes only one camera pose, the uniqueness of such solution is critical to ensure that no equilibrium is lost due to the reconstruction process. Said differently, we want to prove that there exists a one-to-one correspondence between feasible extended features and proper camera poses.

As we will see in the following, in the non-coplanar case, it is possible to show that the (proper) solution Ω computed in Step 3a is unique. Similarly, when the points are coplanar (but non-collinear), one can show that there are exactly one proper and one improper solution in this case. We can then force the reconstructed pose to be the unique proper rotation by using the classical strategy proposed in Step 3b (see, e.g., [56]). Thus, in both the non-coplanar and the coplanar cases, the proposed reconstruction algorithm always finds the unique corresponding camera pose for all solutions of (2.2) and (2.5).

Table 2.3 – Pose reconstruction algorithm.

Input: $\{{}^o\mathbf{a}_i\}_{i=1}^N$ and $\boldsymbol{\xi}$ which is a solution of (2.2)-(2.5)**Output:** $({}^c\mathbf{R}_o, {}^c\mathbf{t}_o)$

1. Determine, through (2.4), $\{{}^c\mathbf{a}_i\}_{i=1}^N$ and define $\mathbf{Q}_i := {}^c\mathbf{a}_i - {}^c\mathbf{k}$ and $\mathbf{P}_i := {}^o\mathbf{a}_i - {}^o\mathbf{k}$, where ${}^c\mathbf{k}$ and ${}^o\mathbf{k}$ represent the points’ centroid in camera- and world-frame, respectively.
 2. Let \mathbf{Q} and \mathbf{P} be the $3 \times N$ matrices that have the corresponding points as columns, define the *covariance matrix* $\mathbf{M} = \mathbf{Q}\mathbf{P}^\top$ and let $\mathbf{M} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^\top$ be its SVD decomposition.
 - 3a. If the points are not coplanar, the best orientation-preserving rotation matrix is $\boldsymbol{\Omega} = \mathbf{U}\mathbf{V}^\top$.
 - 3b. If the points are coplanar, we have instead $\boldsymbol{\Omega} = \mathbf{U} \text{diag}(1, 1, d) \mathbf{V}^\top$, where d is the sign of the determinant of $\mathbf{U}\mathbf{V}^\top$.
 4. Return the reconstructed camera pose as $({}^c\mathbf{R}_o, {}^c\mathbf{t}_o) = (\boldsymbol{\Omega}, {}^c\mathbf{k} - \boldsymbol{\Omega}{}^o\mathbf{k})$.
-

2.2.2 Full solution set of the orthogonal Procrustes’ problem

The orthogonal Procrustes’ problem has been known for a very long time in the literature. The first reference to a solution for it is given by P. H. Schönemann in 1966 [91]. Schönemann proposes a very general discussion on this problem, proving that his solution is always a global minimum of the cost and that it is unique if and only if the covariance matrix $\mathbf{M} = \mathbf{Q}\mathbf{P}^\top$ is full rank. For the case when it is rank deficient, Schönemann simply states that “it appears that [...] the transformation matrix is not unique” [91]. The solution for the orthogonal Procrustes’ problem that is shown in Table 2.3 has been proposed by W. Kabsch in the ’70s. Kabsch seems to be motivated by practical needs, and, even though he briefly discusses the fact that the solution he proposes attains the global minimum of the cost function, he is not concerned with the uniqueness of this solution. In [57], he shows that he is aware that his solution is not uniquely determined in some cases, but he simply provides some additional steps to his procedure to at least find one solution, and he does not investigate the issue further. Two years later, he proposes a slightly different strategy that, only when the singular values of the covariance matrix \mathbf{M} are distinct, selects a rotation which is proper [56]. He also remarks that, seemingly only coincidentally, this new strategy works even when the covariance matrix is of rank 2. In his theoretical development, Kabsch always assumes that \mathbf{M} is full rank, and he considers the loss of rank as an algorithmic, rather than mathematical, issue. In more recent treatments of the orthogonal Procrustes’ problem, like [50] or [49], the uniqueness

of its solution is once again proven when \mathbf{M} is full rank, while the rank deficient case is never discussed.

In the context of visual servoing, it is crucial to assess whether the pose reconstruction leads to a unique possible pose, or if several poses which give the same extended features can exist. Motivated by all of this, we decided to pursue a general study of the orthogonal Procrustes' problem, explicitly deriving its full solution set for any choice of points, as we can see in the following result. We will then see how the uniqueness of the pose reconstruction process in the visual servoing context presented above is an immediate corollary.

Theorem 2.1. *Let $\{\mathbf{a}_i\}_{i=1}^N$ and $\{\mathbf{b}_i\}_{i=1}^N$ be two sets of 3-dimensional points centered on the origin, and let $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{3 \times N}$ be the matrices that have the corresponding points as columns. Then, the orthogonal Procrustes' problem:*

$$\arg \min_{\mathbf{\Omega}^\top \mathbf{\Omega} = \mathbf{I}} \|\mathbf{\Omega} \mathbf{A} - \mathbf{B}\|_F^2, \quad (2.6)$$

can either have 1, 2 or a continuum of solutions depending on the rank of the covariance matrix $\mathbf{M} = \mathbf{B} \mathbf{A}^\top \in \mathbb{R}^{3 \times 3}$. Specifically, if $\mathbf{M} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^\top$ is an SVD decomposition, we have:

- If $\text{rank } \mathbf{M} = 3$, the unique global minimizer is:

$$\mathbf{\Omega} = \mathbf{U} \mathbf{V}^\top. \quad (2.7)$$

- If $\text{rank } \mathbf{M} = 2$, there are two global minimizers, given by:

$$\mathbf{\Omega} = \mathbf{U} \begin{bmatrix} +1 & & \\ & +1 & \\ & & \pm 1 \end{bmatrix} \mathbf{V}^\top. \quad (2.8)$$

- If $\text{rank } \mathbf{M} = 1$, there are two disjoint, 1-dimensional continuum of global minimizers, both of which can be parameterized by $\theta \in [-\pi, \pi)$ as follows, where the determinant of $\mathbf{R}(\theta)$ discriminates between the two sets:

$$\mathbf{\Omega} = \mathbf{U} \begin{bmatrix} 1 & \\ & \mathbf{R}(\theta) \end{bmatrix} \mathbf{V}^\top, \quad \mathbf{R}(\theta) \in \text{O}(2) \quad (2.9)$$

- Finally, if $\text{rank } \mathbf{M} = 0$ (i.e., if \mathbf{M} is the null matrix), then the cost function of (2.6)

does not depend on $\mathbf{\Omega}$, meaning that any $\mathbf{\Omega} \in O(3)$ is a global minimizer.

Proof. Let $\mathbf{M} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$ be an SVD decomposition of $\mathbf{M} = \mathbf{B}\mathbf{A}^\top$. By the properties of the Frobenius product, and assuming that $\mathbf{\Omega}$ is orthonormal, we have that:

$$\begin{aligned} \|\mathbf{\Omega}\mathbf{A} - \mathbf{B}\|_F^2 &= \|\mathbf{\Omega}\mathbf{A}\|_F^2 + \|\mathbf{B}\|_F^2 - \langle \mathbf{\Omega}\mathbf{A}, \mathbf{B} \rangle_F = \|\mathbf{A}\|_F^2 + \|\mathbf{B}\|_F^2 - \langle \mathbf{\Omega}, \mathbf{B}\mathbf{A}^\top \rangle_F = \\ &= \|\mathbf{A}\|_F^2 + \|\mathbf{B}\|_F^2 - \langle \mathbf{U}^\top \mathbf{\Omega} \mathbf{V}, \mathbf{\Sigma} \rangle_F. \end{aligned} \quad (2.10)$$

Since $\mathbf{\Sigma}$ is a diagonal matrix with entries $\sigma_1 \geq \sigma_2 \geq \sigma_3 \geq 0$, and calling $\mathbf{S} = \mathbf{U}^\top \mathbf{\Omega} \mathbf{V}$, we obtain the equivalent formulation:

$$\arg \min_{\mathbf{\Omega}^\top \mathbf{\Omega} = \mathbf{I}} \|\mathbf{\Omega}\mathbf{A} - \mathbf{B}\|_F^2 = \arg \max_{\mathbf{S}^\top \mathbf{S} = \mathbf{I}} \sum_{i=1}^3 s_{i,i} \sigma_i. \quad (2.11)$$

We notice that the number of $\sigma_i \in \mathbb{R}_+$ different than zero is the rank of \mathbf{M} and that, by the orthonormality of \mathbf{S} , it is necessary that all of its entries are between -1 and 1 . Thus, we have the following conditions:

- If all $\sigma_i > 0$, then $\mathbf{S} = \mathbf{I}$ is clearly the unique, global maximizer of (2.11), meaning that $\mathbf{\Omega} = \mathbf{U}\mathbf{V}^\top$.
- If only $\sigma_3 = 0$, then a global maximizer of (2.11) is such that $s_{1,1} = s_{2,2} = 1$. It is easy to verify that the only two orthonormal matrices that satisfy this constraint are such that $s_{3,3} = \pm 1$ with the rest of the entries null, which proves (2.8).
- If only $\sigma_1 > 0$, we similarly have that $s_{1,1} = 1$. Thus, the first row and column of \mathbf{S} must be $(1, 0, 0)$, while its 2×2 bottom-right submatrix just needs to be orthonormal, proving (2.9).
- If \mathbf{M} is the null matrix we trivially get $\|\mathbf{\Omega}\mathbf{A} - \mathbf{B}\|_F^2 = \|\mathbf{A}\|_F^2 + \|\mathbf{B}\|_F^2$, which does not depend on $\mathbf{\Omega}$.

Finally, we remark that the singular value decomposition is not unique when two or more singular values are equal. When this happens, we are allowed to pick any basis of the eigenspace as the corresponding columns of \mathbf{U} and \mathbf{V} . However, this means that we are multiplying \mathbf{U} and \mathbf{V} with an block-orthonormal matrix \mathbf{O} , which, when substituted in the expression of $\mathbf{\Omega}$, either elides (cases 1 and 2) or adds an uninfluential shift to the matrix \mathbf{R} (case 3). \square

Thanks to Theorem 2.1, we have a clear link between the number of solutions of the orthogonal Procrustes problem and the covariance matrix of the two sets of points. We want to remark here that, somewhat oddly, it is not easy to tell whether the covariance

matrix will be rank deficient, and we will later show a few surprising examples of seemingly arbitrary choices of points that admit an infinity of solutions. However, when the two sets of points satisfy the feasibility constraints (2.5), as it is the case in the image-based visual servoing context presented above, it is possible to show that the rank of the covariance matrix is completely determined by points' coplanarity, as we show in the next result.

Corollary 2.1. *Let $\{\mathbf{a}_i\}_{i=1}^N$ and $\{\mathbf{b}_i\}_{i=1}^N$ be two sets of 3-dimensional points centered on the origin and satisfying the feasibility constraints, i.e., they are such that the points' interdistance and the signed volume is the same for both sets. Then, if the points $\{\mathbf{a}_i\}_{i=1}^N$ are either non-coplanar or coplanar (but not collinear), then there is a proper rotation which is the unique solution of (2.6).*

Proof. We already discussed, in the previous section, that the feasibility constraints impose that there exists a proper rotation that maps one set of points into the other, i.e., that there exists a matrix $\mathbf{R} \in \text{SO}(3)$ such that $\mathbf{R}\mathbf{A} = \mathbf{B}$. We want to show here that this matrix is unique.

If the points are non-coplanar, \mathbf{A} is full rank, which implies that $\mathbf{B} = \mathbf{R}\mathbf{A}$ and $\mathbf{M} = \mathbf{B}\mathbf{A}^\top$ are full rank as well. By Theorem 2.1, (2.6) has a unique minimizer, which must coincide with \mathbf{R} being $\|\mathbf{R}\mathbf{A} - \mathbf{B}\|_F^2 = 0$.

On the other hand, if the points are coplanar (but not collinear), \mathbf{A} and $\mathbf{B} = \mathbf{R}\mathbf{A}$ both have rank 2. We argue that $\text{rank } \mathbf{M} = 2$ as well. Since \mathbf{R} is invertible, we have that, for any $\mathbf{x} \in \mathbb{R}^3$:

$$\mathbf{M}\mathbf{x} = \mathbf{R}(\mathbf{A}\mathbf{A}^\top \mathbf{x}) = \mathbf{0} \iff \mathbf{A}\mathbf{A}^\top \mathbf{x} = \mathbf{0}, \quad (2.12)$$

meaning that \mathbf{M} and $\mathbf{A}\mathbf{A}^\top$ have the same kernel and, hence, the same rank. Additionally, it is well known [86] that $\text{rank } \mathbf{A}\mathbf{A}^\top = \text{rank } \mathbf{A}$, and so $\text{rank } \mathbf{M} = 2$. By Theorem 2.1, (2.6) has two global minimizers $\mathbf{\Omega}_1$ and $\mathbf{\Omega}_2$. However, it is clear that $\det \mathbf{\Omega}_1 = -\det \mathbf{\Omega}_2$, meaning that only one of them is a proper rotation, which must coincide with \mathbf{R} for the reasoning above. \square

It should not be surprising that, when the points are all lying on a line, we have a continuum of poses that result in the same extended features: intuitively, any rotation of the camera with respect to the line on which the points lie will not change the extended features. The surprising fact is that, when we do not require that the feasibility constraints are satisfied, we completely lose any relationship between the coplanarity of either cloud of points and the rank of the covariance matrix.

Consider for instance the two following point clouds:

$$\mathbf{A} = \begin{bmatrix} -1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 3 & -1 & -2 \end{bmatrix}, \quad (2.13)$$

$$\mathbf{B} = \begin{bmatrix} -1 & 0 & 0 & 1 \\ 0 & 0 & -1 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}. \quad (2.14)$$

Both clouds lie on a plane that passes through the origin, meaning that $\text{rank } \mathbf{A} = \text{rank } \mathbf{B} = 2$. Their covariance matrix is:

$$\mathbf{M} = \mathbf{B}\mathbf{A}^\top = \begin{bmatrix} 2 & 0 & -2 \\ 1 & 0 & -1 \\ 0 & 0 & 0 \end{bmatrix}, \quad (2.15)$$

which is of rank 1, implying that there is an infinity of solutions. We show the point clouds in Figure 2.1, together with the 1-dimensional set of proper orientations that does not change the error $\|\Omega\mathbf{A} - \mathbf{B}\|_F^2$.

If we consider more points, it is actually easy to construct seemingly arbitrary, non-coplanar point clouds such that the covariance matrix is identically null. Let \mathbf{A} represent a cloud of $N \geq 7$ non-coplanar points centered on the origin. We have that $\dim \ker \mathbf{A} \geq 4$, meaning that we can always find at least three linearly independent vectors in the kernel of \mathbf{A} such that the sum of their elements is null. If we take these three vectors as the rows of \mathbf{B} , we obtain a second cloud of N non-coplanar points such that $\mathbf{B}\mathbf{A}^\top = \mathbf{0}$. We can see such an example in Figure 2.2. This idea can easily be adapted to build similar examples where the ranks of \mathbf{A} , \mathbf{B} and \mathbf{M} can be chosen independently from one another, definitely proving that, outside of the visual servoing context discussed above, no relationship between these values exists in general.

2.3 Verified computation of all critical points

Up to now, we only analyzed the equilibrium condition (2.2)-(2.5) from a purely theoretical point of view, showing that its solutions are always feasible extended features and that we can always reconstruct a unique camera pose which correspond to this solution.

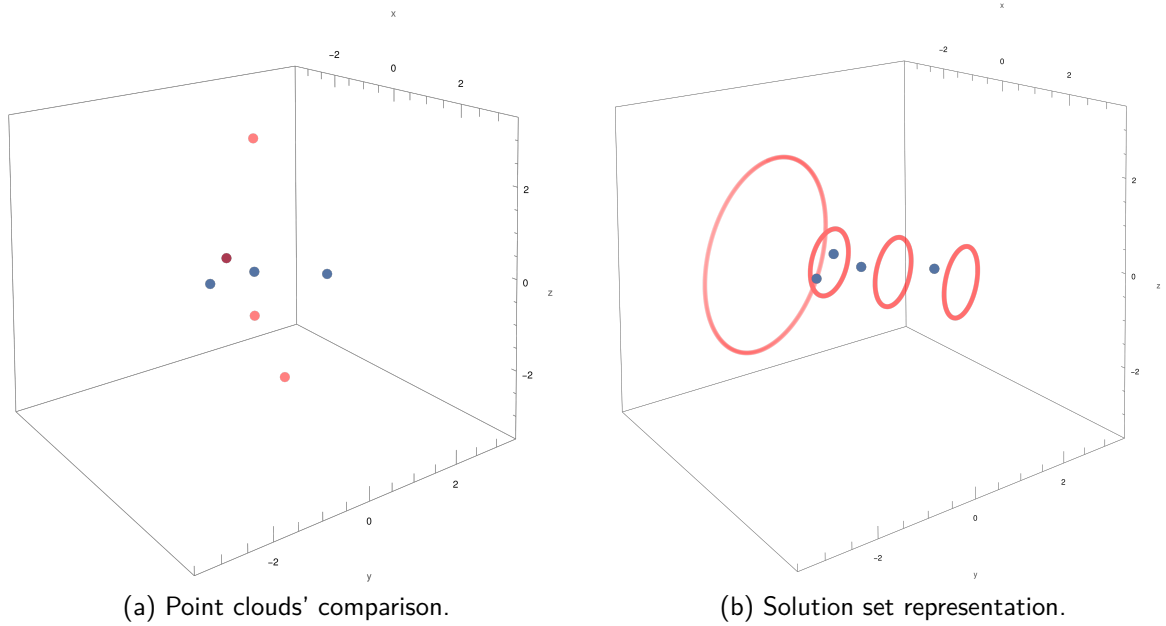


Figure 2.1 – Point clouds (left) and representation of the 1-dimensional set of (proper) solutions (right) of the orthogonal Procrustes problem.

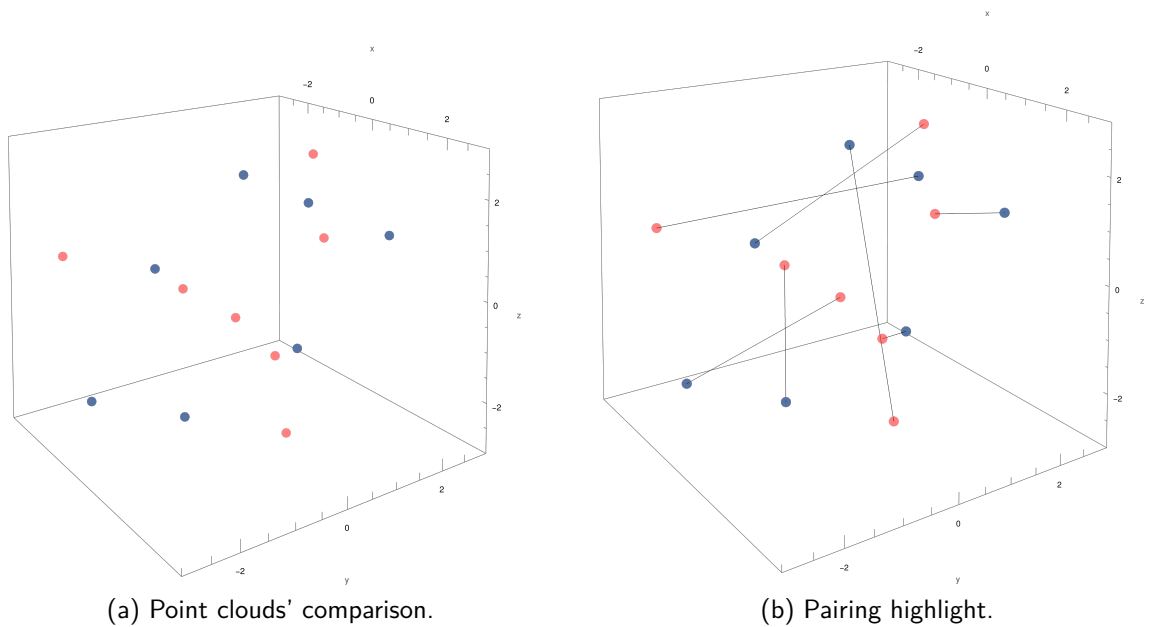


Figure 2.2 – Point clouds (left) with identically null covariance matrix, and the same point clouds with the points pairing highlighted (right).

We now want to find a way to solve the equilibrium condition in practice, analyzing the stability of image-based visual servoing systems using the strategies developed in Chapter 1.

In order to build the connectivity tree for a given system, it is necessary that we know all the critical points of the energy function. This is why, in this section, we propose a method, based on computational algebraic geometry, to find all solutions of a polynomial system of equations in a *verified* way, meaning that this method is guaranteed to not artificially create or miss any of the system’s solutions. We first give a brief presentation of this method and we show that we can solve the equilibrium condition for cases with 4 points, while cases with more points are still out of reach at the current time. We then briefly discuss a modification to the equilibrium condition (2.2)-(2.5) which, in the case when the 4 points are coplanar, significantly improves the solver’s performances.

For the material presented in this section, I want to dearly thank Jorge García Fontán from LIP6 at Sorbonne Université, which allowed us to solve this complicated problem in practice by single-handedly dealing with Gröbner bases methods and `msolve`.

2.3.1 Generalities on computational geometry

This section deals with the resolution of the polynomial systems derived in the previous sections by means of computational algebraic geometry methods and, in particular, Gröbner bases computations.

The tools used to solve polynomial systems fall into two categories: numerical (e.g. Newton’s method, numerical homotopy continuation, interval analysis) [2, 47, 71], and symbolic (e.g. multivariate resultant, cylindrical algebraic decomposition, Gröbner bases) [12, 14]. Numerical methods are generally not successful in finding all solutions in an unbounded domain with certification, see, e.g., a comparison between Gröbner bases and homotopy methods in [45], on problems similar to the ones that we consider. Symbolic methods are thus privileged for applications where an exhaustive search of all solutions is required.

Given a set of polynomials f_1, \dots, f_s in $\mathbb{Q}[\mathbf{x}]$, the ring of polynomials in variables $\mathbf{x} = (x_1, \dots, x_n)$ with rational coefficients, we say that they generate a *polynomial ideal*, i.e. the set defined by all the polynomials that are algebraic combinations of f_1, \dots, f_s :

$$I = \langle f_1, \dots, f_s \rangle = \left\{ \sum_{i=1}^s h_i f_i \mid h_i \in \mathbb{Q}[x_1, \dots, x_n] \right\}. \quad (2.16)$$

For any other polynomial g in the ideal I we have that, if a point $\mathbf{c} \in \mathbb{C}^n$ satisfies $f_1(\mathbf{c}) = \dots = f_s(\mathbf{c}) = 0$, then $g(\mathbf{c}) = 0$ too. This shows that the ideal I is an object to study in order to obtain information about the solutions of the system $f_1(\mathbf{x}) = \dots = f_s(\mathbf{x}) = 0$.

The locus of all the *complex* solutions of this system is called the *algebraic variety* defined by the ideal I , denoted

$$\mathbf{V}(I) = \{\mathbf{c} \in \mathbb{C}^n \mid f(\mathbf{c}) = 0, \forall f \in I\}. \quad (2.17)$$

When the algebraic variety of an ideal consists of finitely many (complex) roots, we say that the ideal is *zero-dimensional* (because its variety is a set of dimension zero). In this case, we define the *degree* of the ideal as the number of complex roots of the system, counted with multiplicity. We refer to [35] for more information about polynomial ideals and varieties.

Gröbner bases are an essential tool in computer algebra for solving problems with polynomial ideals. They can be used, among other things, to determine whether a polynomial is contained in a given ideal, to eliminate a subset of variables from a system of equations, to compute the projection of an algebraic set, or to obtain a rational parametrization of the solutions.

Given a set of polynomials defining an ideal and given a monomial ordering (i.e. a hierarchy of the monomials in the polynomial ring), which must be specified *a priori*, a Gröbner basis is a set of generators of that ideal that have particularly useful computational properties. Different monomial orderings have different effects on the structure of the final Gröbner bases and on the complexity of their computation. For instance, the *lexicographical* ordering provides a triangular description of the ideal, but it is costly to compute, while the *degree reverse lexicographical* ordering results in general in the fastest computations and in polynomials of lower degree. We refer to [12, 35] for a more complete introduction to Gröbner bases and their applications.

The current state-of-the-art methods for computing Gröbner bases are based on the F4 [39] and F5 [40] algorithms by Faugère. The FGLM algorithm [42] for change of ordering allows to compute a lexicographical ordering from a given Gröbner bases with respect to an *easier* monomial ordering. Efficient implementations of these algorithms for solving systems of polynomials exist in computer algebra systems like Magma or Maple, or in libraries such as FGb [41] or `msolve` [9]. Here, we relied on `msolve`, an open-source, high-performance library for computing the real roots of zero-dimensional polynomial

systems, which is based on efficient implementations of F4 and FGLM and on a dedicated univariate real root isolation algorithm.

Using `msolve`, all complex solutions of the system (2.2) and (2.5) are computed together with their multiplicity, real solutions being then easily selected. Note that this system is overconstrained; while this may be a problem for many numerical polynomial solving methods², Gröbner bases computations can deal with it nicely, since the extra relations belonging to the same ideal provide more information about its algebraic structure. Nevertheless, the system (2.2) and (2.5) is still difficult to solve for `msolve`, and we were able to improve the resolution process by breaking a symmetry present in the coplanar case through a change of coordinates invariant to this symmetry.

2.3.2 Symmetry breaking in the coplanar case

In the coplanar case, the solutions of equations (2.2) and (2.5) present a symmetry with respect to a reflection of the scene through the optical center of the camera³. In this section we explain how to exploit this symmetry to our advantage by applying a change of coordinates invariant to this reflection. By means of Gröbner bases and algebraic elimination, we derive a new system of equations in the new variables that is computationally easier to solve.

In particular, consider the transformation that maps the observed points coordinates to their reflection through the camera center by the mapping ${}^c\mathbf{a}_i \mapsto -{}^c\mathbf{a}_i$. Each point with coordinates (\mathbf{s}, \mathbf{Z}) in the space of the extended features is mapped to $(\mathbf{s}, -\mathbf{Z})$ by this reflection. Applying this transformation to the system (2.2) and (2.5) results in a sign change for the polynomials (2.2a-c), while the remaining equations are left unchanged. The solution set is therefore invariant, i.e., if a point $\boldsymbol{\xi} = (\mathbf{s}, \mathbf{Z})$ is a solution, then so is $\boldsymbol{\xi}' = (\mathbf{s}, -\mathbf{Z})$.

Consider now the following change of variables from the Z_i coordinates to a new set of variables θ_{ij} , which is invariant to the reflection described above

$$\mathbf{Z} = (Z_1, \dots, Z_N) \mapsto \boldsymbol{\theta} = (\theta_{12}, \theta_{13}, \dots, \theta_{(N-1)N}), \quad (2.18)$$

with $\theta_{ij} = Z_i Z_j$. By applying this coordinate transformation to (2.2) and (2.5), we expect

2. It is possible to make it well-constrained by taking a suitable subset of equations in (2.5), but it is not detailed here.

3. Note that this is not the case for non-coplanar configurations since this reflection changes the *handedness* of the object, which is preserved by (2.5b).

to remove the spurious symmetric solutions, and derive a system with half the total number of solutions. Applying the change of variables by hand is extremely hard, but it is possible using Gröbner bases and algebraic elimination [35].

From the definition of θ_{ij} we trivially obtain relations between the old and new variables of the form $\theta_{ij} - Z_i Z_j = 0$. For every polynomial f_i in the original system (2.2) and (2.5), we consider the ideal formed by appending these relations, along with a constraint $1 - \ell Z_1 \cdots Z_N = 0$, with ℓ an auxiliary variable, to enforce $Z_i \neq 0$ for all i . This ideal lies in the ring $\mathbb{Q}[\ell, \mathbf{s}, \mathbf{Z}, \boldsymbol{\theta}]$:

$$I_i = \langle f_i, 1 - \ell Z_1 \cdots Z_N, \theta_{12} - Z_1 Z_2, \theta_{13} - Z_1 Z_3, \dots \rangle. \quad (2.19)$$

The ideal I_i is then projected onto the space of variables $(\mathbf{s}, \boldsymbol{\theta})$ by computing a Gröbner basis with respect to an elimination ordering, that is, a monomial ordering which eliminates the variables \mathbf{Z} and ℓ :

$$G_i \leftarrow I_i \cap \mathbb{Q}[\mathbf{s}, \boldsymbol{\theta}]. \quad (2.20)$$

The basis G_i spans all the algebraic combinations of the polynomials in (2.19) that only involve \mathbf{s} and $\boldsymbol{\theta}$. It is equivalent to replacing the variables \mathbf{Z} by $\boldsymbol{\theta}$ in the polynomial f_i . However, the resulting G_i consists not of a single polynomial, but rather of a set of polynomials.

Let us illustrate the algebraic elimination step with an example. Consider the case $N = 4$ and let f_1 be the first polynomial (2.2a). Then, by computing a Gröbner basis for the ideal (2.19) with respect to the elimination ordering we obtain the following six elements, where the first four ones are obtained by multiplying f_1 with Z_i , $i = 1, \dots, 4$:

$$\begin{aligned} G_1 = & [\theta_{12} \theta_{13} (x_4 - x_4^*) + \theta_{12} \theta_{14} (x_3 - x_3^*) + \dots, \\ & \theta_{12} \theta_{23} (x_4 - x_4^*) + \theta_{12} \theta_{24} (x_3 - x_3^*) + \dots, \\ & \theta_{13} \theta_{23} (x_4 - x_4^*) + \theta_{12} \theta_{34} (x_3 - x_3^*) + \dots, \\ & \theta_{12} \theta_{34} (x_4 - x_4^*) + \theta_{14} \theta_{24} (x_3 - x_3^*) + \dots, \\ & -\theta_{12} \theta_{34} + \theta_{14} \theta_{23}, \\ & -\theta_{12} \theta_{34} + \theta_{13} \theta_{24}]. \end{aligned} \quad (2.21)$$

The set union of all the Gröbner bases G_i computed in this form for all the polynomials f_i in (2.2) and (2.5) defines the new system of equations in variables $(\mathbf{s}, \boldsymbol{\theta})$. Once again,

a constraint of the form $1 - \ell \theta_{12} \dots \theta_{1N} = 0$ must be added to remove the solutions with coordinates $\theta_{ij} = 0$ for some i and j , which correspond to solutions with some coordinate $Z_i = 0$. The new system is then

$$\mathbf{G}(\mathbf{s}, \boldsymbol{\theta}) = (G_1, G_2, \dots, G_s, 1 - \ell \theta_{12} \dots \theta_{1N}). \quad (2.22)$$

The system contains more equations and variables than the original one, but the number of complex solutions (i.e. the degree of the ideal) is cut by half, and we observe an improvement in the computation time of up to an order of magnitude (see Section 3.1). For instance, for $N = 4$, the original system contains 13 equations and 12 variables, while the system (2.22) consists of 59 equations in 14 variables.

Every solution $(\mathbf{s}, \boldsymbol{\theta})$ for (2.22) maps to two points (\mathbf{s}, \mathbf{Z}) and $(\mathbf{s}, -\mathbf{Z})$, but we are only interested in those for which all the 3D points are in the semi-space in front of the camera (i.e., $Z_i > 0$ for all i). One can easily verify that a solution with positive Z_i exists if and only if all θ_{ij} are positive. In this case, the depths are recovered by

$$Z_i = \sqrt{\frac{\theta_{ij}\theta_{ik}}{\theta_{jk}}}, \quad (2.23)$$

where j and k are chosen arbitrarily such that $i \neq j \neq k$. The pose reconstruction algorithm in Table 2.3 is then used to determine the proper camera poses corresponding to these solutions.

Overall, the strategy detailed above allow halving the degree of the polynomial ideal by projecting it onto the space of new variables that are symmetry invariant. This leads to a significant reduction in the computation times, resulting in a more effective formulation for computing the critical points of image point-based visual servoing systems.

REPRESENTATIVE TEST CASES IN IMAGE-BASED VISUAL SERVOING

In this chapter, we apply the theory developed throughout Chapters 1 and 2 to a selection of case studies considering configurations of 4 points. For each of the cases, we compute all the equilibria, study their local stability and identify the equilibria's indexes. Then, we simulate the heteroclinic orbits of all unstable equilibria and compute the energy's connectivity tree using Algorithm 1, which is presented in Section 1.3. Additionally, for each example we compute the critical energy levels for which the sublevel sets become non-compact, i.e., the energy level for which the projection onto \mathbb{R}^3 of the sublevel set becomes unbounded and the energy level of the *holes-in-the-wall*, both of which were discussed in Section 1.4. For all cases, we will see that these critical energy levels are significantly higher than the equilibria's energy, which allows to safely talk about the sublevel sets' invariance.




Thanks to the connectivity tree, we can identify an energy level of interest for each test case, which is either the largest controller-independent regions of attraction, in cases where we have multiple minima, or a compact sublevel set when there is a single minimum (i.e., when the desired pose is the only stable equilibrium). We finally compute a paving of the corresponding sublevel set, suitably projected onto \mathbb{R}^3 for visualization purposes.

The chapter is organized as follows. In Section 3.1, we present the test cases that we consider and we give an overview of their analysis. Then, in each subsequent section, we analyze the test cases one by one and present our findings.

All the simulations and plots are made using the `VisualServoingToolbox.jl` Julia package, developed by the author¹.

1. <https://github.com/acolotti/VisualServoingToolbox.jl>

Table 3.1 – Considered test cases and their corresponding equilibria's computation time.

	Configuration	Desired pose		Comp. time	
		ID	${}^o\mathbf{t}_c$		$\theta\mathbf{u}$
1	 ${}^oP = \begin{bmatrix} 0.0 & 1.0 & 0.62 & 0.31 \\ 0.0 & 0.0 & 0.38 & 0.88 \\ 0.0 & 0.0 & 0.0 & 0.31 \end{bmatrix}$	1a	$[-0.66, -0.43, 2.49]$	$[-\pi, 0.0, 0.0]$	22.9h
		1b	$[-3.96, 6.94, -0.41]$	$[1.6, -0.4, 1.2]$	34.7h
2	 ${}^oP = \begin{bmatrix} -0.02 & 0.06 & 0.1 & -0.05 \\ 0.05 & 0.04 & -0.07 & -0.03 \\ 0.0 & 0.0 & 0.0 & 0.0 \end{bmatrix}$	2a	$[-0.11, -0.24, -0.97]$	$[-0.2, 0.2, 1.4]$	25.2h
		2b	$[0.0, 0.0, -1.0]$	$[0.0, 0.0, 0.0]$	2.7h
3	 ${}^oP = \begin{bmatrix} -0.07 & 0.07 & 0.07 & -0.07 \\ 0.07k & 0.07k & -0.07k & -0.07k \\ 0.0 & 0.0 & 0.0 & 0.0 \end{bmatrix}$	$3a, k = \frac{5}{7}$	$[0.0, 0.0, -0.3]$	$[0.0, 0.0, 0.0]$	58.5min
		$3b, k = \frac{6}{7}$	$[0.0, 0.0, -0.3]$	$[0.0, 0.0, 0.0]$	66.0min
		$3c, k = \frac{139}{140}$	$[0.0, 0.0, -0.3]$	$[0.0, 0.0, 0.0]$	103.4min
		$3d, k = 1$	$[0.0, 0.0, -0.3]$	$[0.0, 0.0, 0.0]$	13.1min

3.1 Test cases overview

The test cases that we analyze in this chapter represent classical situations that are encountered in image-based visual servoing. Each considered example is characterized by the configuration of the 4 points in \mathbb{R}^3 together with the desired final pose of the camera. We can split the points configurations into two main categories, which are *coplanar* and *non-coplanar* configurations, and we can split again the coplanar examples by considering desired poses such that the camera optical axis is orthogonal to the points plane or desired poses where this is not the case.

This classification is typical in the visual servoing literature, and several conjectures exist based on it. In particular, for the coplanar case, it is theorized that there are no undesired stable equilibria when the optical axis is orthogonal to the object, while only one undesired stable equilibrium exists when this is not the case. The non-coplanar case is less studied in the literature², and the presence of undesired stable equilibria has been observed only recently [84].

The test cases that we analyze can be found in Table 3.1, where we considered at least one example for each of the categories that we just introduced: we have two examples of non-coplanar configurations (Cases 1a and 1b), one example of coplanar configuration with non-orthogonal desired pose's optical axis (Case 2a), and five examples where the desired pose's optical axis is orthogonal to the object (Cases 2b and 3a to 3d). For each of these examples, we compute all the equilibria and we perform a numerical analysis of their local stability by computing the system's Jacobian in state-space at the equilibrium points and verifying the eigenvalues' signs, where this Jacobian is defined as the derivative of the control input τ_c with respect to the camera pose. Then, by analyzing the unstable equilibria's heteroclinic orbits, we build the corresponding connectivity tree, identify an energy level of interest and trace the (projection onto \mathbb{R}^3 of the) corresponding energy's sublevel set for each of them.

As we discussed in Chapter 1, the transpose, pseudo-inverse and Levenberg-Marquardt controllers give rise to *gradient systems*, meaning that the stability properties of the equilibria are the same for all of them³, regardless of the choice of Levenberg-Marquardt's parameter μ . This equivalence is numerically verified for all the examples below, both by

2. It is important to remark that, in order to have a non-coplanar configuration in practice, the points must lie on some 3-dimensional object. This means that, as the camera moves, one of the points could be covered by the object itself, the reason why these cases are not as studied than the coplanar one.

3. Once again, there is one exception: if one of the equilibria is on a *singularity*, then the pseudo-inverse controller might have different properties than the other two.

the Jacobian’s analysis and by performing extensive simulations with the *pseudo-inverse* and *transpose* controllers. Indeed, the Levenberg-Marquardt controller is highly dependent on the choice of the parameter μ . By changing this parameter, we can, intuitively, adjust the behavior of the controller to have performances more similar to the pseudo-inverse (lower μ) or transpose (higher μ) controllers. Since there is no standard way to tune this parameter, we will discuss only the pseudo-inverse and transpose controllers in the following, considering the Levenberg-Marquardt controller as a compromise between these two.

As anticipated in Chapter 2, the equilibria evaluations were made using `msolve`, a tool developed by the LIP6 at Sorbonne Université, with which we collaborated closely to make the resolution of this problem possible. All computations were performed on a computing server generously made available by the LIP6, which is equipped with an INTEL(R) XEON(R) GOLD 6246R CPU running at 3.40GHz, and with 1.5 TB of total memory. As we anticipated, we only consider cases having $N = 4$ points, since cases with $N > 4$ are too complex to be tackled by the current version of our solver. The computation times can be found in Table 3.1. As expected, exploiting the points’ coplanarity (as discussed in Section 2.3.2) leads to a great reduction in computation time. We can also notice that, as a rule of thumb, the symmetry of the chosen parameters plays a significant role in the resolution complexity, with more symmetric cases requiring much shorter computation times on average.

On the other hand, we computed a guaranteed enclosure of the energy’s sublevel sets using `IBEX`, which is a C++ library for constraint processing over real numbers based on interval arithmetic. We performed all the computations on a mid-range laptop equipped with an INTEL(R) CORE i7 vPRO CPU and with 32 GB of total memory. The computation times were significantly more homogeneous in this case, being all in the 10 hours range. The pavings computed in this way naturally lie in $SE(3)$; we show in the figures below their natural projection⁴ onto \mathbb{R}^3 .

Since `IBEX` relies on an interval version of Newton’s algorithm to rigorously solve the sublevel set’s equation, it is not able to compute a paving for it in a close proximity of a critical point of the energy function, because it represents a *singularity* for the sublevel set’s equation. However, as we saw in Chapter 1, the controller-independent regions of attraction are represented by a sublevel set whose energy level is arbitrarily close to

4. The 3-dimensional surface that we obtain is such that for any position within the surface, there exists at least one corresponding orientation so that the full camera pose is inside the true region of attraction in $SE(3)$.

a saddle point’s one. Thus, only in the cases where we have several minima, we first computed a guaranteed paving of a smaller sublevel set using **IBEX**, and then we “inflated” it up to saddle’s energy level by using a non-verified solver. This procedure can create some artifacts in the paving, as we can see in Figures 3.1c and 3.3c.

Our findings agree with the literature and corroborate the conjectures discussed above, as well as providing for the first time new information regarding the unstable equilibria and the regions of attraction of these systems. The three first test cases Case 1a, Case 1b (four non-coplanar points with two different desired poses) and Case 2a (four planar points with non-parallel desired image plane), whose detailed descriptions are given in Table 3.1, all have one local minimum in addition to the global minimum and two saddles of index 1. They are represented in Figures 3.1, 3.3 and 3.5. These two index 1 saddles have stable manifolds of dimension 5 that form the boundary of the region of attractions of each minimum, as it can be guessed by looking at the corresponding 3-dimensional trajectories (Figures 3.1a, 3.3a and 3.5a). In all these cases, the connectivity tree simply has two leaves, corresponding to the two minima, which get connected at the first saddle that we encounter when we increase the energy level. In each case, two heteroclinic orbits leave the saddles with almost equal initial images and converge toward different minima with their own but similar images, showing similar trajectories in the image (the similitude of the image views of the two heteroclinic orbits is really remarkable in Cases 1a and 2a, as shown in Figures 3.2 and 3.6). However, the corresponding trajectories in $SE(3)$ start in opposite directions and diverge from each other to eventually converge to different minima. Overall, the computation of the minima, saddles and heteroclinic orbits show the similar qualitative dynamics of these three examples. Additionally, we numerically confirm our analysis by performing simulations with initial poses close to the saddles. For each of them, we sample 1000 initial poses uniformly in a small neighborhood of the saddle⁵ and verify how many converge to which stable equilibrium (see Table 3.2 for an overview).

Two cases of planar quadrilaterals with parallel desired image plane are then investigated: Case 2b (four random planar points) and Case 3 (a typical rectangle with centered desired pose). In both cases, these systems have only one minimum and one index 1 saddle, therefore showing a qualitative behavior very different than the previous test cases. The connectivity tree is even simpler in these cases, being all sublevel sets of the energy function connected. Case 3 is investigated for different ratios for the rectangle sides, show-

5. We consider a maximum displacement of 1 cm in position and 0.5° in orientation.

ing that the closer to a square the farther the saddle, see Figures 3.10-3.13. A closed-form formula (3.1) is given for the distance of this saddle with respect to the ratio, proving that it indeed diverges to infinity to finally disappear for a square in Case 3d.

3.2 Test case 1, non-planar configuration

In this example, we consider the case where the 4 points are arranged in a random non-planar configuration. Two different desired poses are considered, one which is such that the optical axis of the camera is perpendicular to one of solid's faces (Case 1a in Table 3.1) while the other is completely arbitrary (Case 1b). In both cases, we have 4 equilibria, one of which coincides with the desired pose (which is stable), a second stable equilibrium (a local minimum) and two unstable equilibria (specifically, two saddles).

It is possible to see the computed equilibria and the corresponding heteroclinic orbits for Cases 1a and 1b in Figures 3.1 and 3.3. Comparing the controllers, we can see (Figures 3.1a-3.3a) that the pseudo-inverse controller's trajectories in $SE(3)$ are slightly different with respect to the transpose controller's ones starting from S_1 . Nevertheless, all trajectories converge to a stable equilibrium, and, as expected, roughly half of them converge to the desired pose, regardless of the chosen control scheme (see Table 3.2). Since, in both cases, the heteroclinic orbits of both saddles are connected to the two minima, then the connectivity trees start with two leaves (the minima) which get connected at the first saddle's energy level, as we can see in Figures 3.1b-3.3b. The controller-independent regions of attraction correspond then to the sublevel sets of the energy function computed at the smallest saddle's energy level, as we can see in Figures 3.1c-3.3c. As expected, due to the difference in energy between the minima, we have that the regions of attraction of the global minima are larger than the local minima ones, even though the asymmetry is more pronounced in Case 1a.

As for the trajectories produced in the image, we can note in Figures 3.2a-3.2c and Figures 3.4a-3.4c that the pseudo-inverse controller produces as expected straight line trajectories for the four image points for converging either to the global minimum or to the local one, while this is not the case for the transpose controller. Interestingly, comparing the global and local minima in image space, we can see that, as it might be expected, they are very similar for Case 1a (Figure 3.9a), while they are surprisingly different for Case 1b (Figure 3.9b).




Figure 3.1 – Computed equilibria, simulations around the saddles and connectivity tree for Case 1a. (a): 3D view, computed equilibria and heteroclinic orbits for the pseudo-inverse (solid) and transpose (dashed) controllers. (b): connectivity tree, with the two critical levels highlighted. (c): controller-independent regions of attraction, computed at S_2 's energy level.

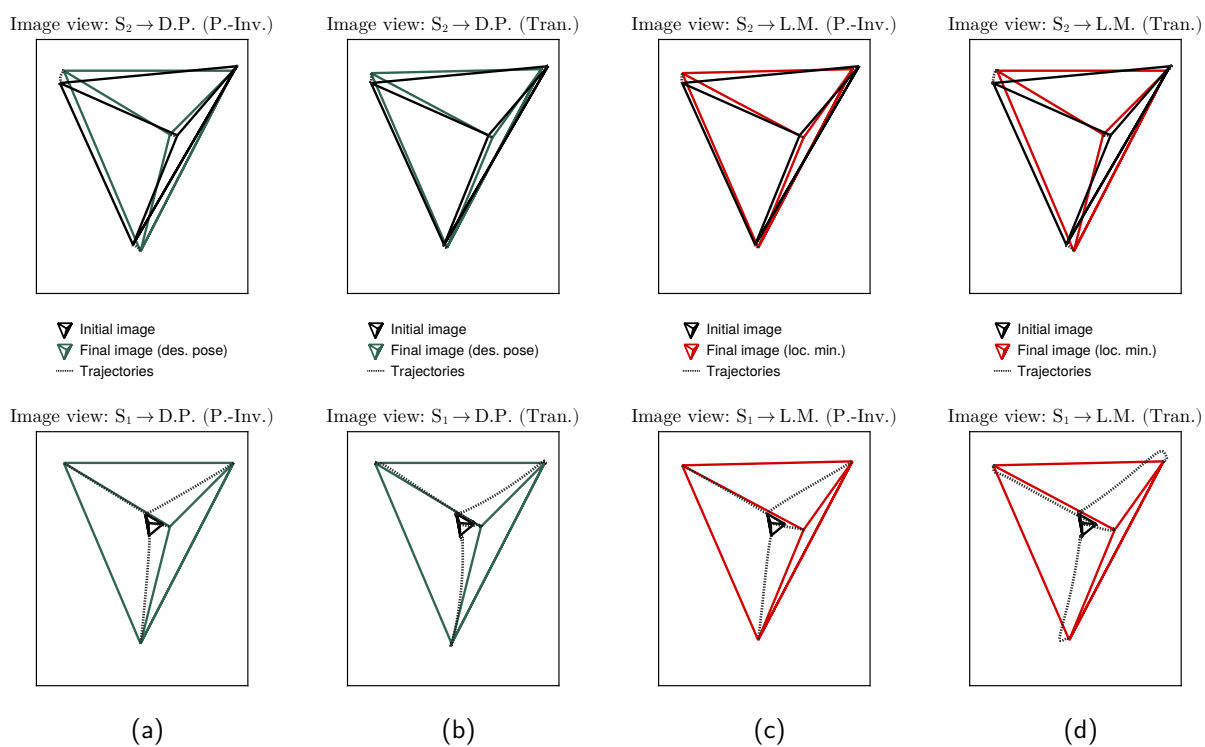


Figure 3.2 – Heteroclinic orbits for Case 1a in image view, for the pseudo-inverse ((a) and (c)) and traspose ((b) and (d)) controllers.




Figure 3.3 – Computed equilibria, simulations around the saddles and connectivity tree for Case 1b. (a): 3D view, computed equilibria and heteroclinic orbits for the pseudo-inverse (solid) and transpose (dashed) controllers. (b): connectivity tree, with the two critical levels highlighted. (c): controller-independent regions of attraction, computed at S_2 's energy level.

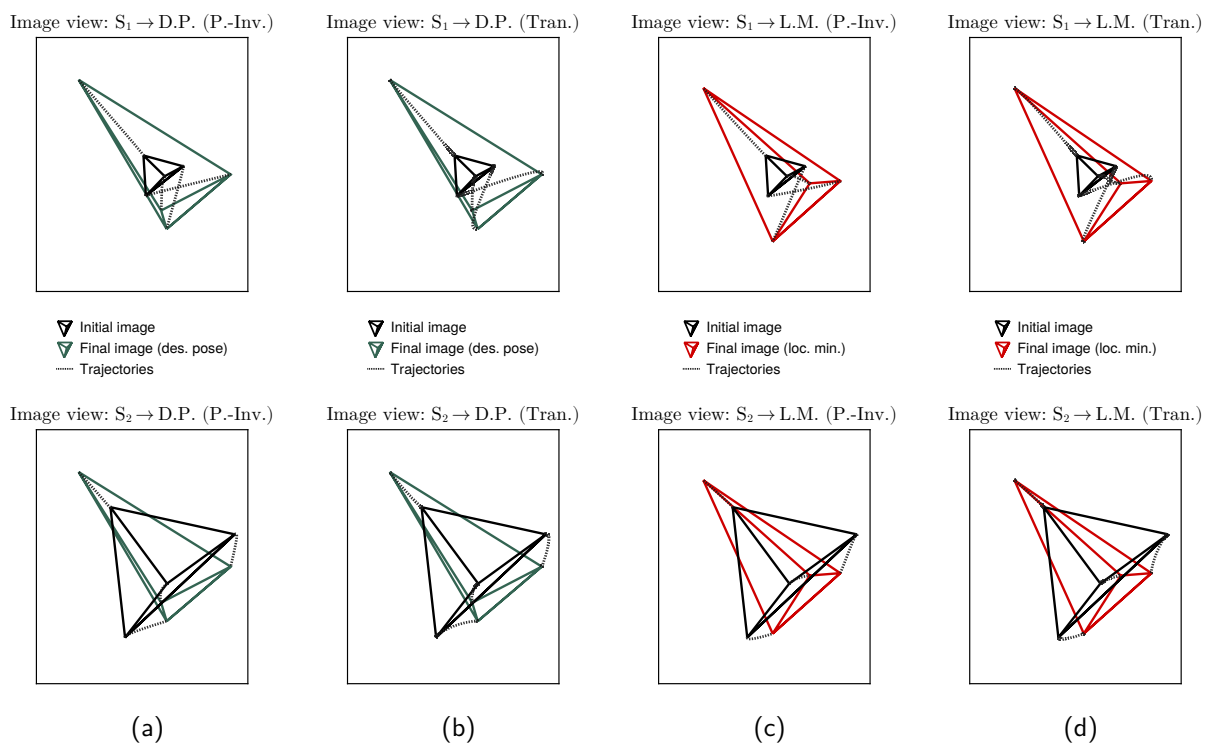


Figure 3.4 – Heteroclinic orbits for Case 1b in image view, for the pseudo-inverse ((a) and (c)) and traspose ((b) and (d)) controllers.

Table 3.2 – Percentage of trajectories (out of 1000 simulations) converging to the desired pose; the remaining ones converge to the other stable equilibrium, when it exists.

		Ex. 1a	Ex. 1b	Ex. 2a	Ex. 2b
P.-Inv.	Sad. 1	51.6%	49.6%	50.1%	100%
	Sad. 2	49%	51.4%	51.6%	N.A.
Transp.	Sad. 1	51.7%	49.4%	51.4%	100%
	Sad. 2	49%	51.7%	52.3%	N.A.

3.3 Test case 2, planar quadrilateral

In this example, we consider the case where the 4 points are coplanar (but non-aligned) and arranged in a random quadrilateral configuration. Once again, two different desired poses are considered, one which is completely arbitrary (Case 2a in Table 3.1) and the other where the quadrilateral is parallel to the image plane (Case 2b). In the first case, we have again 4 equilibria, two of which are stable (being one of them the desired pose), while the other two are saddles. However, in the second case, we only have 2 equilibria, one of which coincides with the desired pose (which is stable), while the other is a saddle. We can note the symmetry between the desired pose and the local minimum with respect to the object plane (even if it is not perfect) and the fact that the saddles are almost parallel to the object plane. These results are fully coherent with the state-of-the-art recalled in the Introduction that a local minimum exists for a planar object in the non-parallel case. In this specific example, we show for the first time that the local minimum is unique and that it does not exist in the parallel case. We conjecture that this might be a general property for this class of cases, even though it still remains an open question.

The computed equilibria for Cases 2a and 2b are illustrated in Figures 3.5 and 3.7, together with trajectories starting in a neighborhood of the saddles and computed using the two controllers. The overall behavior in Case 2a matches the ones in Cases 1a and 1b. In Case 2b, however, we only have one saddle and one stable equilibrium, and all the trajectories converge to the desired pose; interestingly, though, they still do so following two distinct orbits. In turn, this implies that the connectivity tree of Case 2a is analogous to the ones in Cases 1a and 1b, as it is the case for the corresponding controller-independent regions of attraction, both of which are shown in Figures 3.5b-3.5c. On the other hand, the minimum's uniqueness in Case 2b implies that the energy function's sublevel sets always have a single component, meaning that the connectivity tree is simply a straight line




Figure 3.5 – Computed equilibria, simulations around the saddles and connectivity tree for Case 2a. (a): 3D view, computed equilibria and heteroclinic orbits for the pseudo-inverse (solid) and transpose (dashed) controllers. (b): connectivity tree, with the two critical levels highlighted. (c): controller-independent regions of attraction, computed at S_2 's energy level.




Figure 3.6 – Heteroclinic orbits for Case 2a in image view, for the pseudo-inverse ((a) and (c)) and transpose ((b) and (d)) controllers.






Figure 3.7 – Computed equilibria, simulations around the saddle and connectivity tree for Case 2b. (a): 3D view, computed equilibria and heteroclinic orbits for the pseudo-inverse (solid) and transpose (dashed) controllers. (b): connectivity tree, with the two critical levels highlighted. (c): paving of the energy’s level set at $0.1c_S$, with c_S being S ’s energy level.

Image view: $S. \rightarrow D.P. (P.-Inv.)$







 Initial image
 Final image (des. pose)
 Trajectories

Image view: $S. \rightarrow D.P. (Tran.)$








 Initial image
 Final image (des. pose)
 Trajectories

Image view: $S. \rightarrow D.P. (P.-Inv.)$



(a)

Image view: $S. \rightarrow D.P. (Tran.)$



(b)

Figure 3.8 – Heteroclinic orbits for Case 2b in image view, for the pseudo-inverse (a) and transpose (b) controllers.




Figure 3.9 – Comparison between local minimum and desired pose in image space for Cases 1a, 1b and 2a ((a), (b) and (c), respectively).

(Figure 3.7b) and that any sublevel set below c_∞ is a probability-1 controller-independent region of attraction for the minimum. One such sublevel set is shown in Figure 3.7c.

Since, in both cases, the heteroclinic orbits of both saddles are connected to the two minima, then the connectivity trees start with two leaves (the minima) which get connected at the first saddle's energy level, as we can see in Figures 3.1b-3.3b. The controller-independent regions of attraction correspond then to the sublevel sets of the energy function computed at the smallest saddle's energy level, as we can see in Figures 3.1c-3.3c. As expected, due to the difference in energy between the minima, we have that the regions of attraction of the global minima are larger than the local minima ones, even though the asymmetry is more pronounced in Case 1a.

In Case 2a, we can note that the positions of the four points in the image are almost the same for the two stable equilibria and for the saddle located at their middle, while they are very different for the other saddle (as well as for the single saddle in Case 2b). Finally, the percentage of simulated trajectories that converge to the desired pose are given in Table 3.2, where we can see that, as expected, roughly half of the trajectories converge to the local minimum, when it is present.

Note that similar results are obtained when the 4 points are arranged in a rectangle or square configuration, with the exception of the particular case exhibited below.

3.4 Test case 3, from rectangle to square

In this last example, we want to showcase a peculiar behavior obtained while studying the classical case where the 4 points are coplanar and arranged in a rectangle, and where the desired pose is such that the rectangle is centered and parallel to the image plane. We fix the desired pose and we consider four different points configurations, starting with a narrow rectangle (Case 3a in Table 3.1) and making it gradually wider (Cases 3b and 3c) until it becomes a square (Case 3d). For the tested rectangles, there always exist 2 equilibria, one of which coincides with the desired pose (which is stable), while the other is a saddle. As it can be seen in Figures 3.10-3.13, for all three cases the unstable equilibrium is a “mirror image” of the desired pose, i.e., it is rotated 180° around the camera x-axis and it is such that the optical axis is aligned with the desired pose’s optical axis. Furthermore, the trajectories of the four points in the image are no more pure straight line but slightly differ starting from one side of the saddle or the other. We can also note that the four points become aligned in the image when the camera optical center crosses the rectangle plane.

Interestingly, as the rectangle gets closer to a square, the saddle gets pushed farther and farther and, when it becomes a square (Case 3d), we have only one equilibrium, which coincides with the desired pose (which, again, is stable by design). Although, to the best of our knowledge, this is the first time that the uniqueness of the equilibrium is rigorously confirmed for the classical parallel and centered square configuration, it is important to note that this does not imply that the system is globally asymptotically stable; as already said, there exist unstable trajectories for this system, typically for a rotation of 180° around the optical axis [20].

The previous result suggests that the saddle is moved backward at infinity for the square configuration, which motivated us to investigate more deeply this particular case of a parallel centered square. We consider the ratio k between the two sides of the rectangle as a parameter, and we restrict the camera/square poses to a pure translation along the optical axis while fixing their orientation to the saddle’s one. Doing so, the extended features ξ only depend on k and the depth Z w.r.t. the points plane, which significantly simplifies the system (2.2) and (2.5). It is possible to show that this new system is equal to zero if and only if:

$$Z = Z_d \frac{1 + k^2}{|1 - k^2|}, \quad (3.1)$$

where $Z_d > 0$ represents the desired depth w.r.t. the points plane. This means that




Figure 3.10 – Computed equilibria, simulations around the saddle and connectivity tree for Case 3a. (a): 3D view, computed equilibria and heteroclinic orbits for the pseudo-inverse (solid) and transpose (dashed) controllers. (b): connectivity tree, with the two critical levels highlighted. (c): paving of the energy’s level set at $0.07c_S$, with c_S being S ’s energy level.

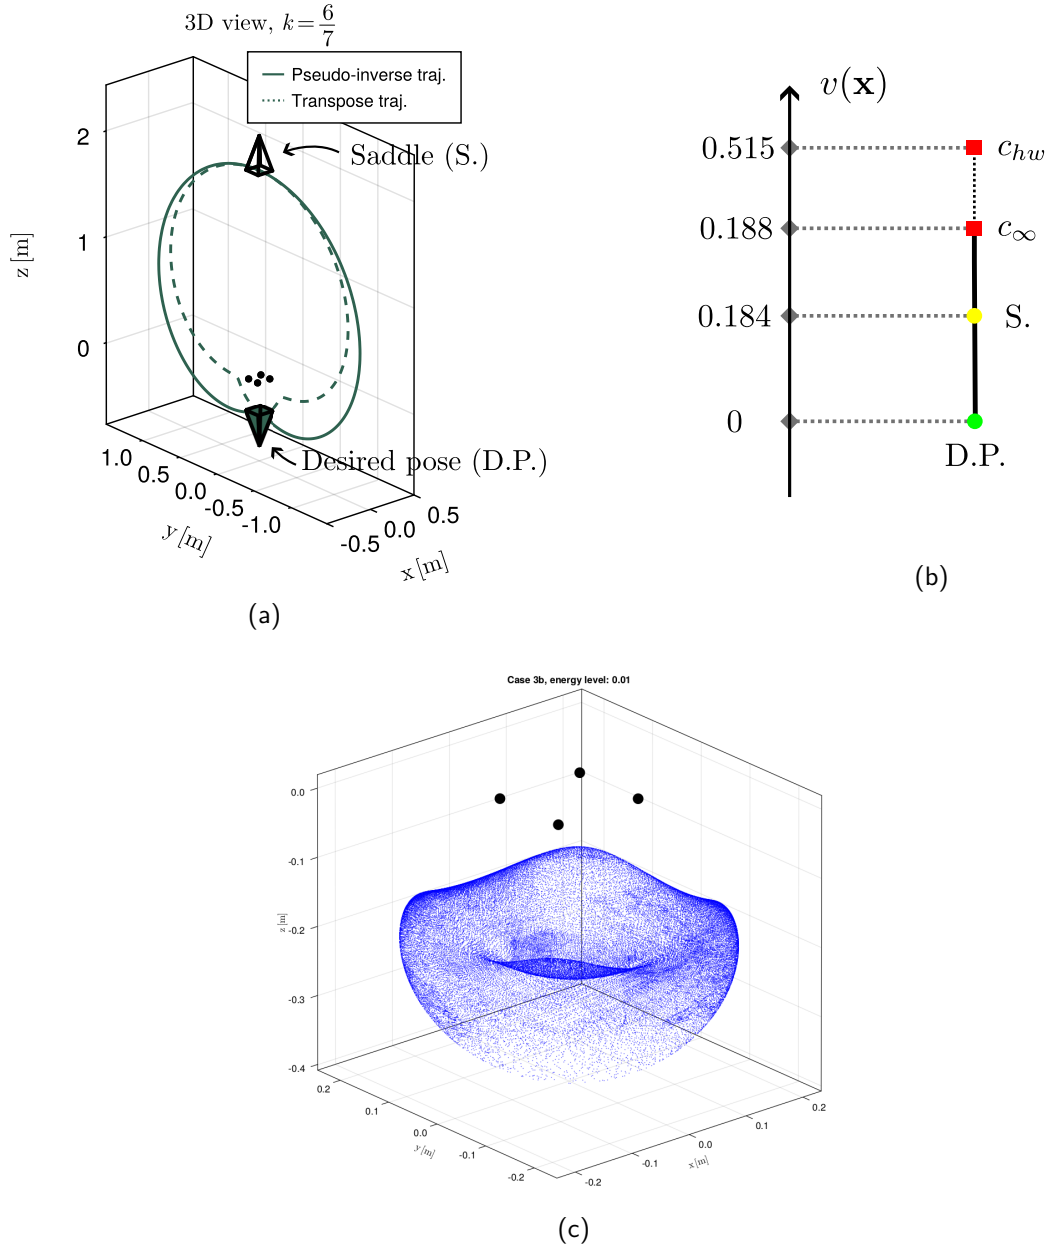


Figure 3.11 – Computed equilibria, simulations around the saddle and connectivity tree for Case 3b. (a): 3D view, computed equilibria and heteroclinic orbits for the pseudo-inverse (solid) and transpose (dashed) controllers. (b): connectivity tree, with the two critical levels highlighted. (c): paving of the energy’s level set at $0.06c_S$, with c_S being S ’s energy level.




Figure 3.12 – Computed equilibria, simulations around the saddle and connectivity tree for Case 3c. (a): 3D view, computed equilibria and heteroclinic orbits for the pseudo-inverse (solid) and transpose (dashed) controllers. (b): connectivity tree, with the two critical levels highlighted. (c): paving of the energy’s level set at $0.05c_S$, with c_S being S ’s energy level.

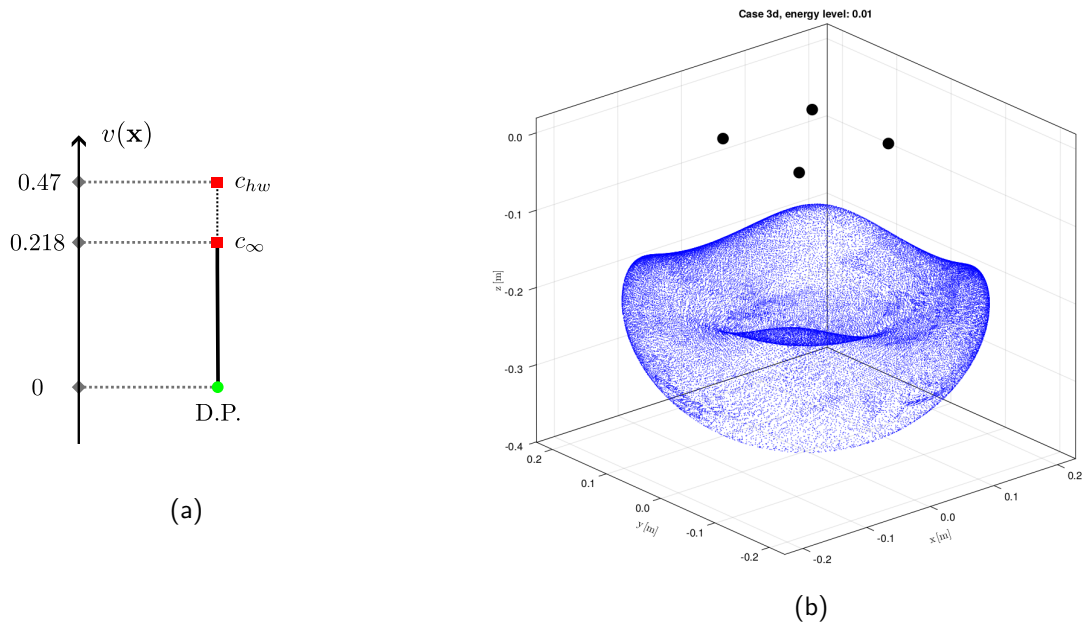


Figure 3.13 – Connectivity tree and region of attraction for Case 3d. (a): 3D view, computed equilibria and heteroclinic orbits for the pseudo-inverse (solid) and transpose (dashed) controllers. (b): connectivity tree, with the two critical levels highlighted. (c): paving of the energy’s level set at $0.4c_\infty$, with c_∞ being the energy level of the first non-compact sublevel set.

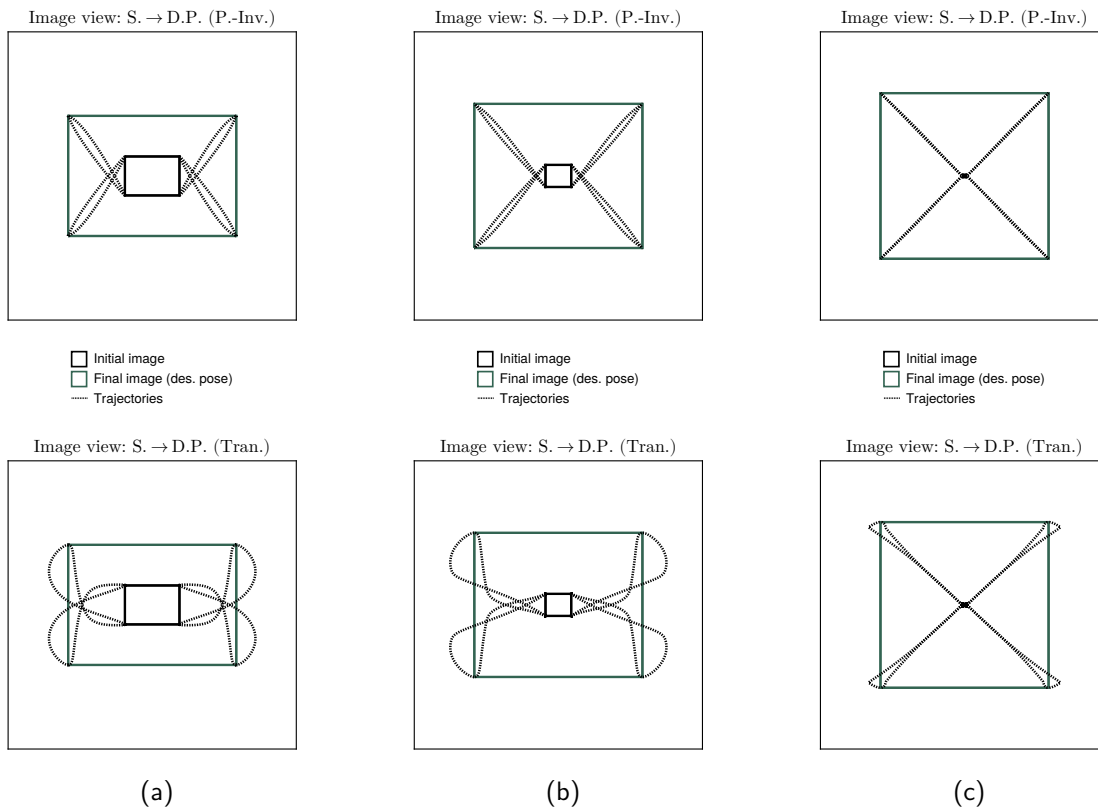


Figure 3.14 – Heteroclinic orbits in image view for Cases 3a-3c (left to right). On the top row: trajectories for the pseudo-inverse controller. On the bottom row: trajectories for the transpose controller.

there is exactly one equilibrium that lies on this restriction for any choice of $k \neq 1$, and it proves that the equilibrium's position moves backwards and tends to infinity as the points get closer and closer to a square configuration. On the other hand, it does not provide any information on the presence of other equilibria lying outside of this restriction, even though we conjecture that no other fixed points exist from the four examples we have considered in this case.

ANALYSIS AND ENERGY-DECREASE ENFORCEMENT FOR OTHER KNOWN CONTROLLERS

In this chapter, we will turn our attention to a different class of well-known image-based visual servoing controllers, that we will call the *desired pose approximation* controllers. While similar to the control laws considered in the previous chapters, these controllers are built using an *estimate* of the interaction matrix, and we will see that, in general, the closed-loop dynamics of these systems are not energy decreasing, at least not for the energy function that naturally arises in image-based visual servoing.

Because of this, we cannot expect their trajectories to be as “well behaved” as their energy-decreasing counterparts. In fact, it often happens that the trajectories of these systems *fail* (i.e., the Cartesian coordinates in the screen run off to infinity) in finite time, especially if we take initial poses which are not in the immediate proximity of the desired pose. However, thanks to the new theoretical tools that we derived in Chapter 2, we are able to show that, in several common situations, these controllers only have one stable equilibrium. This motivates us to try to modify the desired pose approximation controllers, forcing them to be energy decreasing, in an attempt to combine the advantages of both energy-decreasing and non-energy-decreasing strategies. We develop a method, that we call *meta-controller*, which allows to modify any controller so that the closed-loop system is energy decreasing, and we apply it to the desired pose approximation controllers, obtaining a substantial improvement in the regions of attraction’s size with respect to the classical energy-decreasing image-based visual servoing controllers.

The chapter is organized as follows. In Section 4.1, we present the desired pose approximation controllers and show that the equilibrium condition derived in Chapter 2 can be easily adapted to exhaustively compute all of their equilibria. Additionally, in analogy to the energy-decreasing controllers that we already encountered, we show they all

have the same equilibria. In Section 4.2, we consider a set of representative test cases for these controllers, showing both the presence of failed trajectories and the stable equilibrium's uniqueness. Then, in Section 4.3, we present our meta-controller and we show that, under shallow hypotheses, it is always continuous. Finally, we test it on two examples, performing a qualitative comparison between the desired pose's region of attraction for the meta-controller and the classical pseudo-inverse controller to show the improvement of the region's size.

4.1 Desired pose approximation controllers

We consider in this section three new image-based visual servoing controllers, which represent an approximated version of the three energy-decreasing controllers that we considered up to now, namely the transpose, pseudo-inverse and Levenberg-Marquardt controllers. Analogously to what we did back in Chapter 1, all of these new control laws will be based on the general form:

$$\boldsymbol{\tau}_c = -\lambda \mathbf{C}_{\mathbf{e}^*} \mathbf{e} \text{ with } \lambda \in \mathbb{R}_+, \quad (4.1)$$

with three possible choices for $\mathbf{C}_{\mathbf{e}^*}$, namely:

- *Transpose*: $\mathbf{C}_{\mathbf{e}^*} = \mathbf{L}_{\mathbf{e}^*}^\top$,
- *Pseudo-inverse*: $\mathbf{C}_{\mathbf{e}^*} = \mathbf{L}_{\mathbf{e}^*}^+$,
- *Levenberg-Marquardt*: $\mathbf{C}_{\mathbf{e}^*} = \left(\mathbf{L}_{\mathbf{e}^*}^\top \mathbf{L}_{\mathbf{e}^*} + \mu \mathbf{I} \right)^{-1} \mathbf{L}_{\mathbf{e}^*}^\top$, with $\mu \in \mathbb{R}_+$,

where $\mathbf{L}_{\mathbf{e}^*}$ represents the usual interaction matrix evaluated at the desired extended features, i.e. $\mathbf{s} = \mathbf{s}^*$ and $\mathbf{Z} = \mathbf{Z}^*$. We call these controllers the *desired pose approximation* controllers.

Considering an approximation of the interaction matrix in the control loop is typical in applications, since, in general, we cannot measure the points' depths Z_i directly [54]. The simple approximation that we consider here is interesting because we can almost use as-is the equilibrium condition (2.2)-(2.5) that we derived in Chapter 2 to compute the equilibria of the desired pose approximation controllers. In order to derive the new equilibrium condition, though, we want to understand whether these three controllers have the same equilibria or not. Similarly to the energy-decreasing case, we have the following lemma, which shows how the equilibria characterization for all strategies can be actually reduced to the case of the transpose controller.

Lemma 4.1. *Given the approximated interaction matrix $\mathbf{L}_{\mathbf{e}^*}$, the system's equilibrium configurations, characterized by:*

$$\mathbf{L}_{\mathbf{e}^*}^\top \mathbf{e} = 0, \quad (4.2)$$

are the same for the pseudo-inverse, transpose and Levenberg-Marquardt controllers.

Proof. We argue that $\tau_c = 0$ if and only if $\mathbf{L}_{\mathbf{e}^*}^\top \mathbf{e} = 0$ for all three controllers. This is obvious for the transpose controller. It is well known [13] that $\ker \mathbf{L}_{\mathbf{e}^*}^+ = \ker \mathbf{L}_{\mathbf{e}^*}^\top$, proving the equivalence for the pseudo-inverse controller. Finally, by the positive semi-definiteness of $\mathbf{L}_{\mathbf{e}^*}^\top \mathbf{L}_{\mathbf{e}^*}$, $(\mathbf{L}_{\mathbf{e}^*}^\top \mathbf{L}_{\mathbf{e}^*} + \mu \mathbf{I})$ is invertible for all $\mu > 0$, meaning that $(\mathbf{L}_{\mathbf{e}^*}^\top \mathbf{L}_{\mathbf{e}^*} + \mu \mathbf{I})^{-1} \mathbf{L}_{\mathbf{e}^*}^\top \mathbf{e} = 0$ if and only if $\mathbf{L}_{\mathbf{e}^*}^\top \mathbf{e} = 0$, concluding the proof. \square

Thus, not only the desired pose approximation controllers have the same equilibria, but the equilibrium condition is the same as the corresponding energy-decreasing one, with the exception of using the approximated interaction matrix $\mathbf{L}_{\mathbf{e}^*}$ in the place of the true one. The desired pose approximation equilibrium condition can be found in Table 4.1. When put together with the feasibility constraints (2.5), we obtain a polynomial system of equations with a significantly smaller degree with respect to the energy-decreasing equilibrium condition, meaning that we can once again solve it with the computational algebraic geometry methods described in Section 2.3. Once all solutions are computed, it is then possible to retrieve the camera pose using the reconstruction strategy presented in Section 2.2.

4.2 Representative test cases for desired pose approximation

Similarly to the energy-decreasing case, we propose here a set of representative test cases. We study three examples of points' configurations, analyzing a representative set of desired poses for each example. Table 4.2 summarizes the examples' parameters, where the computation times can be found as well. As expected, exploiting the points' coplanarity (as discussed in Section 2.3.2) leads to a great reduction in computation time. We can also notice that, as a rule of thumb, the symmetry of the chosen parameters plays a significant role in the resolution complexity, with more symmetric cases requiring much shorter computation times on average. Moreover, we can see the significant increase of time as we add new points to the configuration.

Table 4.1 – Explicit expression for the desired pose approximation’s equilibrium condition.

$$\sum_{i=1}^N \left[\left(\prod_{j \neq i}^N Z_j^* \right) (x_i^* - x_i) \right] = 0 \quad (4.3a)$$

$$\sum_{i=1}^N \left[\left(\prod_{j \neq i}^N Z_j^* \right) (y_i^* - y_i) \right] = 0 \quad (4.3b)$$




$$\sum_{i=1}^N \left[\left(\prod_{j \neq i}^N Z_j^* \right) \left(x_i^*(x_i - x_i^*) + y_i^*(y_i - y_i^*) \right) \right] = 0 \quad (4.3c)$$

$$\sum_{i=1}^N x_i^* y_i^* (x_i - x_i^*) + (1 + (y_i^*)^2) (y_i - y_i^*) = 0 \quad (4.3d)$$

$$\sum_{i=1}^N (1 + (x_i^*)^2) (x_i^* - x_i) + x_i^* y_i^* (y_i^* - y_i) = 0 \quad (4.3e)$$

$$\sum_{i=1}^N (x_i y_i^* - y_i x_i^*) = 0. \quad (4.3f)$$

Table 4.2 – Desired pose approximation test cases and their corresponding equilibria’s computation time.

	Configuration	Desired pose		Comp. time	
		ID	${}^o\mathbf{t}_c$		$\theta\mathbf{u}$
1	 ${}^oP = \begin{bmatrix} 0.0 & 1.0 & 0.62 & 0.31 \\ 0.0 & 0.0 & 0.38 & 0.88 \\ 0.0 & 0.0 & 0.0 & 0.31 \end{bmatrix}$	1a	$\begin{bmatrix} -0.66, -0.43, 2.49 \end{bmatrix}$	$\begin{bmatrix} -\pi, 0.0, 0.0 \end{bmatrix}$	27.3s
		1b	$\begin{bmatrix} -3.96, 6.94, -0.41 \end{bmatrix}$	$\begin{bmatrix} 1.6, -0.4, 1.2 \end{bmatrix}$	116s
2	 ${}^oP = \begin{bmatrix} 0.0 & 1.0 & 0.62 & 0.31 & -0.31 \\ 0.0 & 0.0 & 0.38 & 0.88 & 0.5 \\ 0.0 & 0.0 & 0.0 & 0.31 & -0.31 \end{bmatrix}$	2a	$\begin{bmatrix} -0.66, -0.43, 2.49 \end{bmatrix}$	$\begin{bmatrix} -\pi, 0.0, 0.0 \end{bmatrix}$	254.8s
		2b	$\begin{bmatrix} -3.96, 6.94, -0.41 \end{bmatrix}$	$\begin{bmatrix} 1.6, -0.4, 1.2 \end{bmatrix}$	2.2h
3	 ${}^oP = \begin{bmatrix} -0.07 & 0.07 & 0.07 & 0.07 & -0.07 \\ 0.07 & 0.07 & -0.07 & -0.07 & -0.07 \\ 0.0 & 0.0 & 0.0 & 0.0 & 0.0 \end{bmatrix}$	3a	$\begin{bmatrix} 0.0, 0.0, -0.5 \end{bmatrix}$	$\begin{bmatrix} 0.0, 0.0, 0.0 \end{bmatrix}$	0.07s
		3b	$\begin{bmatrix} 0.0, 0.0, -0.3 \end{bmatrix}$	$\begin{bmatrix} 0.0, 0.0, 0.0 \end{bmatrix}$	0.07s
		3c	$\begin{bmatrix} 0.0, 0.0, -0.14\sqrt{2} \end{bmatrix}$	$\begin{bmatrix} 0.0, 0.0, 0.0 \end{bmatrix}$	0.08s
		3d	$\begin{bmatrix} 0.0, 0.0, -0.1 \end{bmatrix}$	$\begin{bmatrix} 0.0, 0.0, 0.0 \end{bmatrix}$	0.07s

As we already discussed, this system is not energy decreasing anymore. In fact, we will see that, in spite of a simpler expression due to the usage of a fixed interaction matrix, it has in general more complex dynamics. Fixed points are not related anymore to the error function, with the obvious exception of the global minimum, that still corresponds to an attractive fixed point. Unstable fixed points can still be associated to an index corresponding to the dimension of their unstable manifold, but we cannot use the information coming from the heteroclinic orbits to construct the regions of attraction anymore. Surprisingly, all tested random cases (four and five points non-planar configurations, with two desired poses in each case) have exactly one global minimum and one saddle. Even more surprisingly, the saddle now has index 2, meaning that a surface of heteroclinic trajectories leaves the saddle, here all converging toward the desired pose. All of them have arbitrarily similar initial images and identical final images, but run very different trajectories both in $SE(3)$ and in image space, see Fig. 4.1 and Figures 4.3-4.4. Together with some trajectories that do not converge due to physical constraints, this quantitative analysis shows a complex global dynamics of this controller.

4.2.1 Test case 1, 4 points non-planar configuration

In this example, we reanalyze Case 1 from Chapter 3, where 4 points are arranged in a random non-planar configuration. The same desired poses are considered, one which is such that the optical axis of the camera is perpendicular to one of solid's faces (Case 1a in Table 4.2) while the other is completely arbitrary (Case 1b). Contrarily to the energy-decreasing examples, we now have 2 equilibria for both cases, one of which coincides with the desired pose (which is stable) and one unstable equilibrium (specifically, a saddle).

These equilibria are illustrated in Figures 4.1-4.2, where we also show a selection of trajectories, both in camera state- and image space, starting in a close proximity of the saddle. Comparing these trajectories with the ones obtained in the previous chapter, we have a striking qualitative difference between the trajectories' behavior in the two settings. For both desired poses, the trajectories seem to develop on a surface connecting the saddle to the desired pose (even though the surfaces appear to be significantly more regular for Case 1b than Case 1a). This behavior is linked to the dimension of the saddle's unstable manifold: checking the eigenvalues of the system's Jacobian computed at the saddle, we can see that in both cases we have two positive eigenvalues.

Moreover, with respect to the energy-decreasing controllers, we have a new behavior that arises, which is the presence of *failed* trajectories, i.e., trajectories that come too close

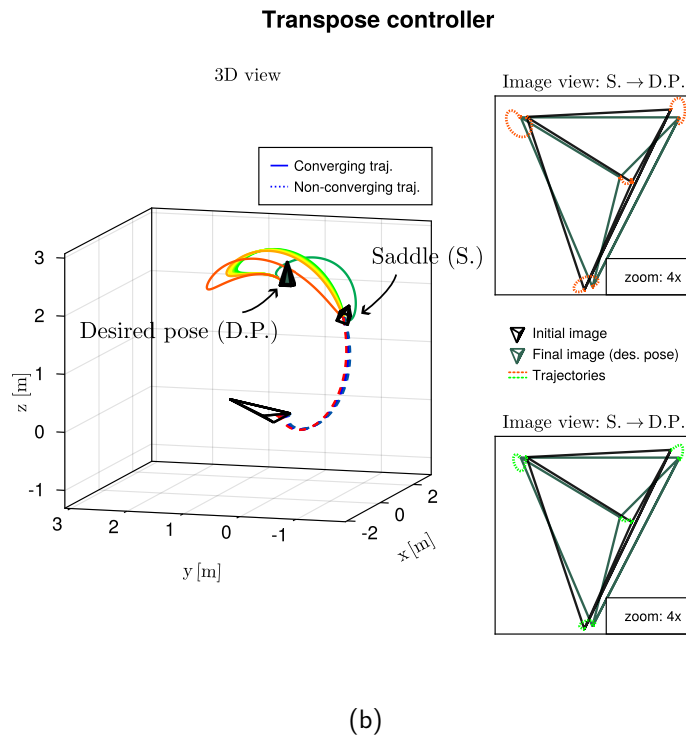



Figure 4.1 – Computed equilibria and simulations around the saddle for Case 1a. (a): computed equilibria and a selection of heteroclinic orbits, both in 3D and image space, for the pseudo-inverse controller. (b): idem for the transpose controller.

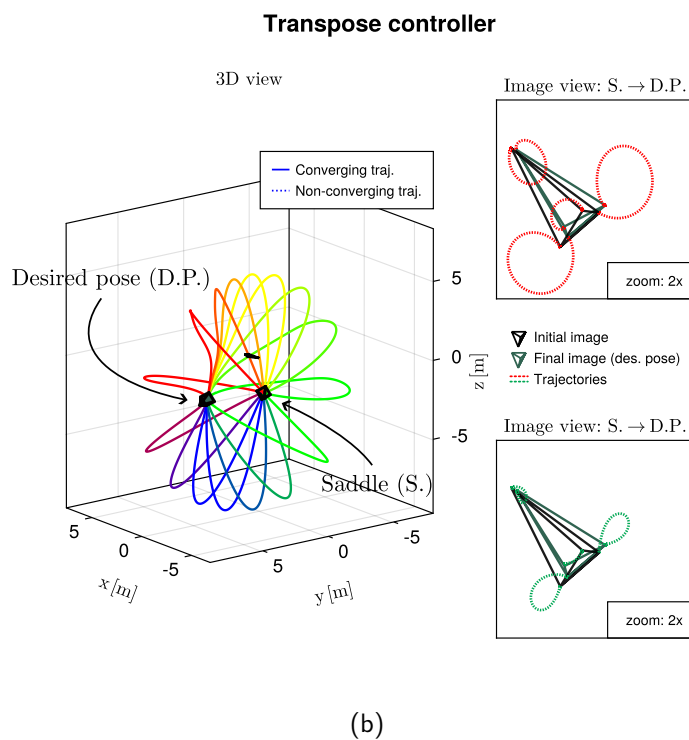
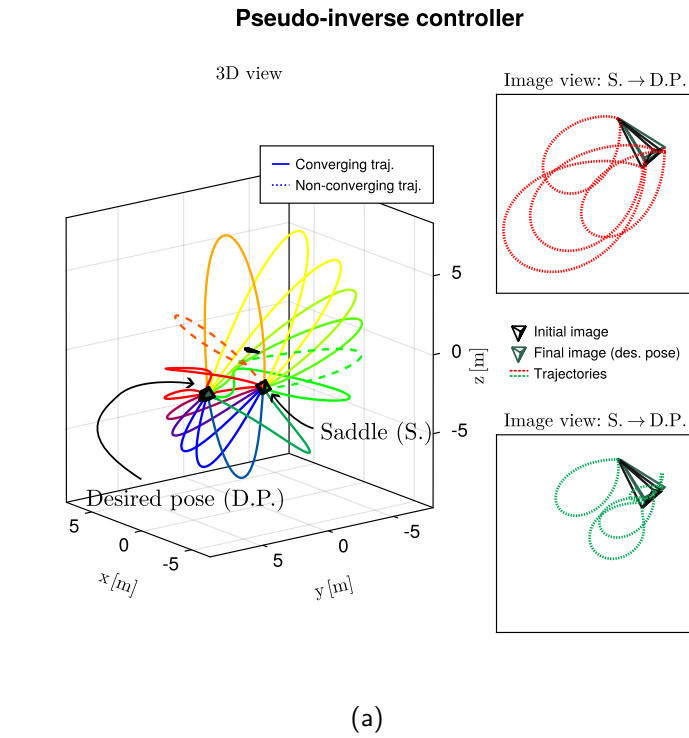


Figure 4.2 – Computed equilibria and simulations around the saddle for Case 1b. (a): computed equilibria and a selection of heteroclinic orbits, both in 3D and image space, for the pseudo-inverse controller. (b): idem for the transpose controller.

to the object so that the depth Z of at least one object point becomes null (shown with dashed lines in the figures). This phenomenon, however, is not surprising: the employed control laws are based on a constant and very coarse approximation of the true interaction matrix, and, by design, we expect them to be effective only when the camera is reasonably close to the desired pose. In general, we have no guarantees for trajectories with arbitrary initial conditions, and, as we will see, this behavior arises in almost all the examples considered in this section. Additionally, considering the trajectories in the image, we see that the points do not move at all along straight lines (while it is the case for the energy-decreasing controllers). Once again, this validates the approximated nature of this class of controllers, which leads to more unpredictable (but, also, richer) behavior.

Similar examples where the 4 points are arranged in a coplanar rectangle configuration lead to comparable results with respect to Case 1, the reason why these cases were left out from the current discussion.

4.2.2 Test case 2, 5 points non-planar configuration

In this example, we study a variation of Cases 1*a* and 1*b*. We consider a configuration of 5 points, formed by the 4 points considered in the previous example and an additional, random point, and we use the very same desired poses (Cases 2*a* and 2*b* in Table 4.2). Again, in both cases, we have 2 equilibria, a stable one (coinciding with the desired pose) and a saddle. Even though we have significant similarities with the previous example, quite surprisingly the saddles do not bear any evident resemblance to what we found above.

In Figures 4.3-4.4, we can see the computed equilibria and a selection of trajectories for Cases 2*a* and 2*b*. While in Case 2*b* the trajectories seem to develop on a surface connecting the saddle to the desired pose, matching the behavior of Cases 1*a* and 1*b*, in Case 2*a* the trajectories quickly converge to one single orbit, as it was the case for the saddles seen in Chapter 3. Once again, this behavior is due to the dimension of the saddle's unstable manifold: in Case 2*a* we have an index 1 saddle, while there are two positive eigenvalues in Case 2*b*.

Finally, we can see that the trajectories in the image are far from developing along straight lines, and they present an erratic behavior that changes conspicuously depending on the initial pose, which showcases, once again, the approximated nature of these control laws.





Figure 4.3 – Computed equilibria and simulations around the saddle for Case 2a. (a): computed equilibria and heteroclinic orbits, both in 3D and image space, for the pseudo-inverse controller. (b): idem for the transpose controller.

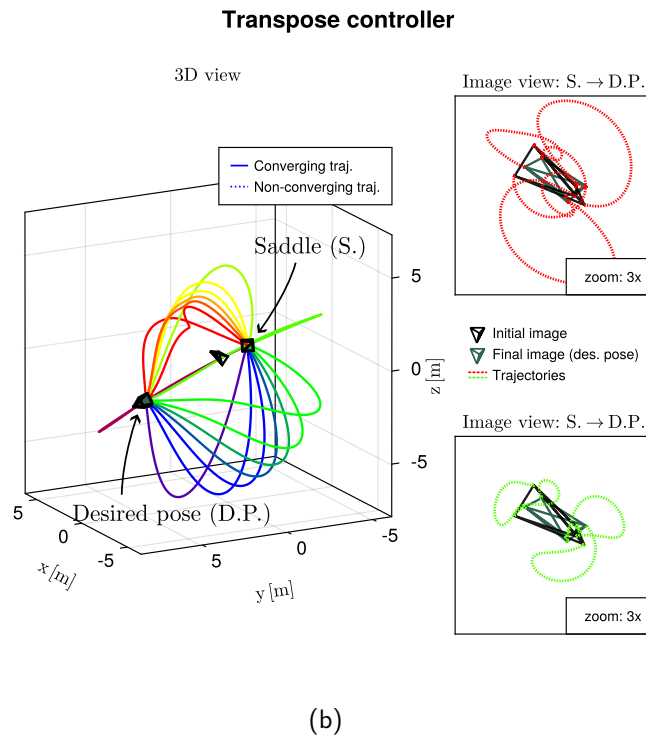



Figure 4.4 – Computed equilibria and simulations around the saddle for Case 2b. (a): computed equilibria and a selection of heteroclinic orbits, both in 3D and image space, for the pseudo-inverse controller. (b): idem for the transpose controller.




Figure 4.5 – Cases 3a-3c (left to right). On the top row: 3D view, computed equilibria and heteroclinic orbits for the pseudo-inverse (solid) and transpose (dashed) controllers. On the bottom row: image view, one of the heteroclinic orbits for the pseudo-inverse controller.

4.2.3 Test case 3, planar square

In this last example, we want again to showcase a curious behavior that arises in the classical case where the 4 points are coplanar and arranged in a square, and where the desired pose is such that the square is centered and parallel to the image plane. We consider four different desired poses, all with the same orientation and lying on the same axis, with only the distance from the points plane changing: we start farther away (Case 3a in Table 4.2) and we get increasingly closer (Cases 3b to 3d) to the points plane. In the first three cases, we have 5 equilibria, four of which are saddles, while the fifth one is the desired pose (stable by design). All the saddles lie on the points plane, and they are arranged symmetrically with respect to the points. In the fourth case, instead, we only have 1 equilibrium, coinciding with the desired pose. By looking at Fig. 4.5, we can see

that, as the desired pose approaches the points plane, the four saddles get closer to the four points (Cases 3a and 3b), until they reach a position where they are superposed to the points (Case 3c). Drawing the desired pose even closer, the saddles eventually disappear (Case 3d), or, more precisely, they become such that at least one point is behind the camera optical center, which excludes them to be a feasible solution.

In all cases, the four saddles share the same relative orientation, and only their distance to the points changes from case to case, which motivated us to investigate this example more deeply. In the same spirit of the analysis carried out in Section 3.4, we consider the desired depth Z_d as a parameter, and we restrict the camera's position to be on one of the axes where the saddles' lie¹ and we fix its orientation to the saddle's one. With this restriction, the extended features ξ depend only on the (signed) distance d w.r.t. the closest point, with $d > 0$ when the point lies in front of the camera (i.e., when the point's depth in camera-frame is positive). Substituting into the system (4.3)-(2.5), it is possible to show that the equilibrium condition is equal to zero if and only if:

$$d = \frac{1}{2} (Z_d - \sqrt{2}L), \quad (4.4)$$

where $L > 0$ is the length of the square's side. This means that these four, symmetrical saddles are present for every choice of $Z_d \geq \sqrt{2}L$, and, as Z_d decreases, the saddles positions get closer and closer to the points, until they eventually become superposed to them at the equality. Interestingly, these saddles are solutions of (4.3)-(2.5) even in the case where $d < 0$, and they tend to converge to the center of the square as Z_d tends to 0; however, these solutions lose physical meaning, since one of the points is behind the camera.

4.3 Meta-controller to enforce energy decrease of an arbitrary control law

In this section, we will present a way to force an arbitrary control law to be energy decreasing for a given energy function. The idea behind our approach is quite straightforward: we want to use the information coming from the energy function's gradient to define a set of energy-decreasing velocities, so that any velocity in the set is such that

1. We can choose any of the four axes; it has no impact on the result due to the symmetry of the problem.

$\mathcal{L}_f v(\mathbf{x}) < 0$, for any \mathbf{x} which is not a critical point of the energy function. Then, we take the velocity computed with the initial, non-energy-decreasing controller and project it on this newly defined set of *admissible* velocities. In this way, we create a new, energy-decreasing control law which is as close as possible to the original controller, effectively combining the advantages of both approaches.

While keeping the image-based visual servoing applications in mind, we argue that this control strategy can be applied in robotics at large. In fact, throughout the section, we will assume that we have a fully actuated system of the form $\dot{\mathbf{x}}(t) = \mathbf{u}(t)$, where $\mathbf{u}(t)$ is an n -dimensional input and the state-space \mathcal{M} is an n -dimensional manifold. This is the model that we considered up to now for visual servoing free-floating camera systems, but it is also a classical model for robotic arms, holonomic mobile robots and other fully actuated robotic systems [59], as long as the dynamical effects are negligible. In practice, in a real robotics system, we typically have embedded low-level controllers that manage the force-velocity conversion, meaning that, at least as long as we assume that the movements are slow enough, the robot will be able to precisely follow the velocity provided in input.

Additionally, we consider an output $\mathbf{y} = \mathbf{s}(\mathbf{x}) \in \mathbb{R}^m$, with $m \geq n$, depending smoothly on \mathbf{x} but possibly nonlinear. We assume that the system is observable so that the state can be accurately reconstructed from the output. Once again, this assumption is typical in (but not restricted to) visual servoing where, e.g., the depth of the observed points need to be estimated to compute the interaction matrix, and the impact of uncertainties state estimation as well as the impact of imperfect state estimation on the stability are investigated independently. As a consequence, the closed loop control law $\mathbf{u}(\mathbf{y})$ can actually depend on the state and the closed loop system is eventually $\dot{\mathbf{x}} = \mathbf{u}(\mathbf{x})$.

Being given a desired output \mathbf{y}^* , which usually comes from a desired state \mathbf{x}^* , i.e., $\mathbf{y}^* = \mathbf{s}(\mathbf{x}^*)$, one naturally defines an energy function

$$v(\mathbf{x}) = \frac{1}{2} \|\mathbf{e}(\mathbf{x})\|^2, \quad (4.5)$$

with $\mathbf{e}(\mathbf{x}) = \mathbf{s}(\mathbf{x}) - \mathbf{y}^*$. As usual, the energy function v can be highly nonlinear with possibly several local minimizers. The gradient of (4.5) can be defined by $\nabla v(\mathbf{x}) = \mathbf{J}(\mathbf{x})^\top \mathbf{e}(\mathbf{x})$, where $\mathbf{J}(\mathbf{x})$ is the Jacobian of the error such as $\dot{\mathbf{e}} = \mathbf{J}(\mathbf{x})\dot{\mathbf{x}}$, for the canonical scalar product associated to the vector representation of tangent vectors.

With these assumptions, we can generalize the energy-decreasing control laws that we encountered in image-based visual servoing for any system of the form $\dot{\mathbf{x}}(t) = \mathbf{u}(t)$. They

are easily defined using the gradient information of $v(\mathbf{x})$: the transpose, pseudo-inverse and Levenberg-Marquardt controllers are given by $\mathbf{u} = -\lambda \mathbf{J}(\mathbf{x})^\top \mathbf{e}(\mathbf{x})$, $\mathbf{u} = -\lambda \mathbf{J}(\mathbf{x})^+ \mathbf{e}(\mathbf{x})$ and $\mathbf{u} = -\lambda (\mathbf{J}(\mathbf{x})^\top \mathbf{J}(\mathbf{x}) + \mu \mathbf{I})^{-1} \mathbf{J}(\mathbf{x})^\top \mathbf{e}(\mathbf{x})$, respectively. Analogously, by taking $\mathbf{J}(\mathbf{x}) = \mathbf{J}(\mathbf{x}^*)$, we can define the corresponding desired pose control law.


The section is organized as follows. We propose a general construction that builds the composite energy-decreasing law closest to a given non-energy-decreasing law, which is introduced on an illustrative example in Section 4.3.1. In Section 4.3.2, this construction is generalized to arbitrary systems of the form $\dot{\mathbf{x}}(t) = \mathbf{u}(\mathbf{x})$, and an explicit algorithm for the computation of the meta-controller is provided. In Section 4.3.3, we prove that the meta-controller is continuous under reasonable hypotheses on the energy function and the non-energy-decreasing control law. Finally, the framework is applied to visual servoing in Section 4.4, using the desired pose approximation pseudo-inverse controller as the non-energy-decreasing law. Additionally, we propose in this section a methodology to assess the region of attraction of visual servoing system, by projecting error sublevel sets in \mathbb{R}^3 and deciding a preferred starting orientation. We illustrate this on two test cases taken from Chapter 3, and we show that the region of attraction of the global minimizer is greatly enlarged using the newly obtained control law.

4.3.1 Illustrative example


We consider the system $\dot{\mathbf{x}} = \mathbf{u}$ where $\mathbf{x}(t), \mathbf{u}(t) \in \mathbb{R}^2$, and the output $\mathbf{y}(t) \in \mathbb{R}^3$ whose components are the squared distances $y_i = \|\mathbf{x} - \mathbf{a}_i\|^2$ to three known points \mathbf{a}_i . We consider $\mathbf{a}_1 = (5, 3)$, $\mathbf{a}_2 = (-5, 3)$ and $\mathbf{a}_3 = (5, 2)$, and desired outputs $\mathbf{y}^* = (34, 34, 29)$ so that the origin $\mathbf{x}^* = 0$ is the global minimizer of the energy function $v(\mathbf{x}) = \|\mathbf{e}(\mathbf{x})\|^2$, where $\mathbf{e}(\mathbf{x}) = \mathbf{y}(\mathbf{x}) - \mathbf{y}^*$. This energy function also has a local minimizer $\mathbf{x}_{\text{loc}} \approx (0.24, 5.34)$ and a saddle $\mathbf{x}_{\text{sad}} \approx (-0.12, 2.90)$. The Jacobian matrix of the system is

$$\mathbf{J}(\mathbf{x}) = \frac{d\mathbf{y}}{d\mathbf{x}} = 2 \begin{pmatrix} -5 + x_1 & -3 + x_2 \\ 5 + x_1 & -3 + x_2 \\ -5 + x_1 & -2 + x_2 \end{pmatrix}. \quad (4.6)$$


The phase portraits of the transpose controller $\mathbf{u} = -\mathbf{J}(\mathbf{x})^\top \mathbf{e}(\mathbf{x})$ and the pseudo-inverse one $\mathbf{u} = -\mathbf{J}(\mathbf{x})^+ \mathbf{e}(\mathbf{x})$ are shown in Figure 4.6a and Figure 4.6b respectively. As expected, the fixed points of both systems correspond to the critical points of the energy function. Also as expected, the transpose controller shows some directions with rapid and slow con-




(a) Phase portrait of the transpose controller.



(b) Phase portrait of the pseudo-inverse controller.



(c) Phase portrait of the composite law $\dot{\mathbf{x}} = \mathbf{u}_p(\mathbf{x})$.



(d) Zoom, dashed areas are $U(\mathbf{x}_1)$ and $U(\mathbf{x}_2)$.

Figure 4.6 – Stream plots of the three energy-decreasing controllers. Red lines represent the stable manifold of the saddle, that separate the regions of attraction of the two minimizers, while the blue lobes represent the sublevel sets of $v(\mathbf{x})$ at the saddle’s energy level. The last graphics shows one point entering the sublevel set and another sliding on it.

vergence near the global minimizer, while the pseudo-inverse one shows a homogeneous convergence, since they correspond to the steepest descent and Gauss-Newton optimization algorithms, respectively. The derivative of the systems vector fields evaluated at the saddle have the same eigenvectors associated to stable and unstable directions, hence the stable manifolds of the saddle have the same initial direction. In both cases, these stable manifolds separate the regions of attractions of the minimizers. All of these observations agree with the energy-decreasing systems’ analysis carried out in Chapter 1.

We want now to consider a control law $\dot{\mathbf{x}} = \mathbf{u}_i(\mathbf{x})$ that is not energy decreasing globally, but has appealing properties. To emphasize the benefit of the meta-controller approach, we choose for this simple example $\mathbf{u}_i(\mathbf{x}) = -100\mathbf{x}$. This control law is trivially globally asymptotically stable, and critical points of the energy function are not anymore fixed point of the system. We specify requirements on the composite law $\mathbf{u}_p(\mathbf{x})$ to be defined: first, we require it enforces a strong enough error decrease with respect to the steepest descent:

$$\nabla v(\mathbf{x})^\top \mathbf{u}_p(\mathbf{x}) \leq -\epsilon \nabla v(\mathbf{x})^\top \nabla v(\mathbf{x}), \quad (4.7)$$

where $\epsilon > 0$ is a fixed ratio, here $\epsilon = 0.01$. The second requirement is technical, to enforce the continuity of the composite law: we require that its norm is not too much larger than the norm of the gradient:

$$\|\mathbf{u}_p(\mathbf{x})\| \leq \rho \|\nabla v(\mathbf{x})\|, \quad (4.8)$$

where $\rho > 0$ is a fixed ratio, here $\rho = 100$. These two constraints correspond to a set $U(\mathbf{x})$ of admissible commands:

$$U(\mathbf{x}) = \{\mathbf{u} \in \mathbb{R}^n : \frac{1}{\rho^2} \mathbf{u}^\top \mathbf{u} \leq \mathbf{g}^\top \mathbf{g} \leq -\frac{1}{\epsilon} \mathbf{g}^\top \mathbf{u}\}, \quad (4.9)$$

where $\mathbf{g} = \nabla v(\mathbf{x})$ for typesetting, so that the two requirements (4.7) and (4.8) express equivalently $\mathbf{u}_p(\mathbf{x}) \in U(\mathbf{x})$. The set $U(\mathbf{x})$ is the intersection of a ball and a half space, which is non-empty if and only if $\epsilon \leq \rho$. We require that $\epsilon < \rho$ so that $U(\mathbf{x})$ has a non-empty interior whenever $\nabla v(\mathbf{x}) \neq 0$ (check that $\mathbf{u} = -\frac{\epsilon+\rho}{2} \nabla v(\mathbf{x})$ satisfies strictly both constraints). Otherwise we have $U(\mathbf{x}) = \{0\}$. We are now in position to define the composite law $\mathbf{u}_p(\mathbf{x})$: since we want $\mathbf{u}_p(\mathbf{x})$ to be as much similar to $\mathbf{u}_i(\mathbf{x})$ as possible, it is natural to define $\mathbf{u}_p(\mathbf{x})$ as the projection of $\mathbf{u}_i(\mathbf{x})$ onto $U(\mathbf{x})$:

$$\mathbf{u}_p(\mathbf{x}) = \arg \min_{\mathbf{u} \in U(\mathbf{x})} \|\mathbf{u} - \mathbf{u}_i(\mathbf{x})\|^2. \quad (4.10)$$

The set $U(\mathbf{x})$ being nonempty and convex, and the quadratic objective function strictly convex, there is one unique minimal argument and $\mathbf{u}_p(\mathbf{x})$ is well defined.

The phase portrait of $\dot{\mathbf{x}} = \mathbf{u}_p(\mathbf{x})$ is shown in Figure 4.6c. We see that the global minimizer has a region of attraction that is greatly enlarged with respect to the transpose and pseudo-inverse control laws. This is possible thanks to the very untypical saddle stable manifold, which appears to be broken (and, thus, it is not a smooth manifold). Because of this, similarly to an example that we saw back in Section 1.2.2, this system

does not satisfy the hypotheses of Theorem 1.3, entailing that the composite control law is not differentiable at the saddle. Interestingly, the broken stable manifold of the saddle seems to slide on the saddle level set, which is a specific feature of this non-differentiable composite control law. To understand this property, note first that since the control law is energy decreasing, any trajectory starting inside the upper half sublevel containing the local minimizer will converge to this local minimizer. Now, if it starts outside but close to sublevel set then two cases arise: if it starts at \mathbf{x}_1 where $\mathbf{u}_i(\mathbf{x}_1)$ is strongly energy decreasing, see Figure 4.6d, then $\mathbf{u}_p(\mathbf{x}_1) = \mathbf{u}_i(\mathbf{x}_1)$ and the trajectory eventually converges to the local minimizer. If it starts at \mathbf{x}_2 where $\mathbf{u}_i(\mathbf{x}_2)$ is strongly energy-increasing instead, see again Figure 4.6d, then the set of energy-decreasing velocities is approximately parallel to the sublevel set and projecting onto it gives rise to a velocity approximately tangent to the sublevel set and therefore the trajectory approximately slides on it.

The non-differentiability of the composite control law at the saddle opens up questions related to its continuity properties in general, which will be discussed later in the section. Nevertheless, if the saddle converges to the global minimizer for the input control law then sliding on the saddle sublevel set corresponds to the largest region of attraction of any energy-decreasing control law locally in the vicinity of the saddle. The global impact of the composite control law on the region of attraction needs to be investigated for each system. The next section presents the detailed application of the process to a general sensor-based system, built on the pseudo-inverse of the desired pose approximation's Jacobian matrix, which is well known to have good properties but is not globally energy decreasing.

4.3.2 General definition of the meta-controller

Definitions (4.10) and (4.9) hold in the general n -dimensional case as well. For a fixed \mathbf{x} , finding the control law $\mathbf{u}_p(\mathbf{x})$ amounts to solving a constrained nonlinear optimization problem. This optimization problem involves a convex quadratic cost function, a linear constraint and a convex quadratic constraint, and can be solved formally using Karush-Kuhn-Tucker optimality conditions. This lengthy computation process is not detailed here, but instead a graphical illustration of the structure of the problem and its formal solution is presented.

Due to the axial symmetry of the feasible set $U(\mathbf{x})$ around the vector $\nabla v(\mathbf{x})$, the projection $\mathbf{u}_p(\mathbf{x})$ of $\mathbf{u}_i(\mathbf{x})$ onto $U(\mathbf{x})$ lies in the plane formed of the vectors $\nabla v(\mathbf{x})$ and $\mathbf{u}_i(\mathbf{x})$. For a fixed $\nabla v(\mathbf{x})$ proportional to $(1, 1)$, Figure 4.7 presents the feasible set and the projection process emphasizing its piecewise structure. In green, the feasible set $U(\mathbf{x})$ is

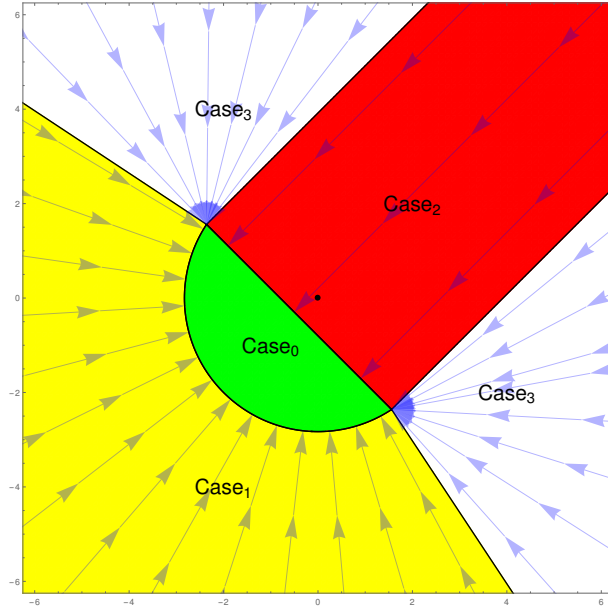


Figure 4.7 – Projection of \mathbf{u}_i onto the set of feasible commands (depicted in green). Cases 0 to 3 represent the regions where different constraints are active.

the intersection of a disk and a half-plane. Blue arrows show how different \mathbf{u}_i projects onto $U(\mathbf{x})$. One can obviously see that four cases arise depending on the position of the vector \mathbf{u}_i , which are represented with four colors. Each four cases enjoys a simple expression, which is presented in Algorithm 2.

4.3.3 Continuity of the meta-controller

We investigate the continuity of the composite control law $\mathbf{u}_p(\mathbf{x})$ in more general settings than the specific form introduced in (4.9). As usual, we assume that $v(\mathbf{x}) \in \mathcal{C}^2$ with nondegenerate critical points. Additionally, we assume that the set of feasible commands $U(\mathbf{x})$ is defined by a family of continuous and convex inequality constraints, and we show that $U(\mathbf{x})$ being convex bounded with nonempty interior when $\nabla v(\mathbf{x}) \neq 0$ and being “well behaved” at critical points of $v(\mathbf{x})$ are sufficient conditions for the continuity of the composite law.

In order to provide a formal expression for these hypotheses, we need to introduce a few key concepts from the theory of set-valued functions, in particular their upper and lower *hemicontinuity* and the celebrated Berge Maximum Theorem [3, Theorem 17.31], which provides sufficient conditions for characterizing continuity of the (set-valued) minimal

Data: $\mathbf{x}, \epsilon, \rho$
Result: Solution \mathbf{u}_p of (4.10) and (4.9)
 $\mathbf{u}_i, \mathbf{g}, r \leftarrow \mathbf{u}_i(\mathbf{x}), \nabla v(\mathbf{x}), \rho \|\nabla v(\mathbf{x})\|$
 $\mathbf{p}_2 \leftarrow \mathbf{u}_i - \frac{\mathbf{g}^\top \mathbf{u}_i}{\mathbf{g}^\top \mathbf{g}} \mathbf{g} - \epsilon \mathbf{g}$
if $\|\mathbf{u}_i\|^2 \leq r^2$ **and** $\mathbf{g}^\top (\epsilon \mathbf{g} + \mathbf{u}_i) \leq 0$ **then**
 $\mathbf{u}_p \leftarrow \mathbf{u}_i$; // Case₀
else if $\|\mathbf{u}_i\|^2 > r^2$ **and** $\mathbf{g}^\top (\epsilon \|\mathbf{u}_i\| \mathbf{g} + r \mathbf{u}_i) \leq 0$ **then**
 $\mathbf{u}_p \leftarrow \frac{r \mathbf{u}_i}{\|\mathbf{u}_i\|}$; // Case₁
else if $\|\mathbf{p}_2\|^2 \leq r^2$ **then**
 $\mathbf{u}_p \leftarrow \mathbf{p}_2$; // Case₂
else
 $\mathbf{u}_p \leftarrow \frac{(\mathbf{p}_2 + \epsilon \mathbf{g}) \sqrt{r^2 - \|\epsilon \mathbf{g}\|^2}}{\|\mathbf{p}_2 + \epsilon \mathbf{g}\|} - \epsilon \mathbf{g}$; // Case₃
end

Algorithm 2: Analytic solution of (4.10) and (4.9).

argument of a nonlinear parametric optimization problem. We only provide here some basic definitions, due to [3]², and we refer the interested reader to [3, 6] for a general introduction on set-valued functions.

Definition 4.1. Let \mathcal{X} and \mathcal{Y} be topological spaces. A **set-valued function** U from \mathcal{X} to \mathcal{Y} , denoted with $U : \mathcal{X} \rightrightarrows \mathcal{Y}$, assigns to each \mathbf{x} in \mathcal{X} a subset $U(\mathbf{x})$ of \mathcal{Y} . Additionally, we say that U is:

- **Upper hemicontinuous (uhc)** at \mathbf{x} if for every neighborhood Y of $U(\mathbf{x})$, there is a neighborhood X of \mathbf{x} such that $\mathbf{z} \in X$ implies $U(\mathbf{z}) \subset Y$,
- **Lower hemicontinuous (lhc)** at \mathbf{x} if for every open set Y such that $U(\mathbf{x}) \cap Y \neq \emptyset$, there is a neighborhood X of \mathbf{x} such that $\mathbf{z} \in X$ implies $U(\mathbf{z}) \cap Y \neq \emptyset$,
- **Continuous** at \mathbf{x} if it is both upper and lower hemicontinuous.

As with functions, we say that U is uhc/lhc/continuous if it is so at every point of \mathcal{X} .

We can assess the continuity of $\mathbf{u}_p(\mathbf{x})$ by applying Berge Maximum Theorem to the parametric optimization problem (4.10), which requires, in addition to some shallow conditions on the objective function $\|\mathbf{u} - \mathbf{u}_i(\mathbf{x})\|^2$ which are easily verified, that the set-valued

2. The definitions of upper and lower hemicontinuity are not universally agreed upon, some authors requiring additional properties for the upper hemicontinuity (see, e.g., [75] for an overview).

function $U(\mathbf{x})$ is continuous. In the following proposition, we provide sufficient conditions for the continuity of $U(\mathbf{x})$.

Proposition 4.1. *Let $U : E \rightrightarrows \mathbb{R}^n$ be defined as $U(\mathbf{x}) = (V \circ \nabla v)(\mathbf{x})$, where $\nabla v(\mathbf{x})$ is continuous and $V : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ is defined as $V(\mathbf{g}) = \{\mathbf{u} \in \mathbb{R}^n : c_i(\mathbf{g}, \mathbf{u}) \leq 0\}$. Assume that:*

- (H1) *The constraints $c_i(\mathbf{g}, \mathbf{u})$ are continuous with respect to \mathbf{g} and \mathbf{u} and convex with respect to \mathbf{u} ,*
- (H2) *$\mathbf{g} \neq 0$ implies that $V(\mathbf{g})$ is bounded and $\text{int } V(\mathbf{g}) \neq \emptyset$,*
- (H3) *For all $\epsilon > 0$, there exists $\delta > 0$ such that $\|\mathbf{g}\| \leq \delta$ implies $V(\mathbf{g}) \subseteq B_\epsilon(0)$.*

Then, $U(\mathbf{x})$ is continuous.

Remark. *The hypothesis (H3) encodes the continuity of the set-valued map $V(\mathbf{g})$ at $\mathbf{g} = 0$, which entail $V(0) = \{0\}$. This technical hypothesis is actually easy to check in practice, see the proof of Corollary 4.1.*

Proof. Since the composition of continuous set-valued functions is continuous [3, Theorem 17.23] and being $\nabla v(\mathbf{x})$ continuous by hypothesis, we just need to prove that $V(\mathbf{g})$ is continuous. When $\mathbf{g} \neq 0$, (H1) and (H2) imply the continuity of $V(\mathbf{g})$ by [7, Theorem 3.2.1]; in particular, $\text{int } V(\mathbf{g}) \neq \emptyset$ trivially implies that the *characteristic index set* (see [7]) of $V(\mathbf{g})$ is the empty set. On the other hand, (H3) implies that $V(0) = \{0\}$, and it is trivial to verify that both the definitions of upper and lower hemicontinuity are satisfied at $\mathbf{g} = 0$. Thus, $V(\mathbf{g})$ is continuous on \mathbb{R}^n . \square

We can now show that the composite control law \mathbf{u}_p is a continuous function when the parametric optimization problem (4.10) satisfies the hypotheses of Proposition 4.1. Finally, as a concluding corollary, we show that the set of feasible commands (4.9) lead to a continuous control law.

Proposition 4.2. *Consider the composite control law \mathbf{u}_p defined in (4.10) and let \mathbf{u}_i be continuous and $U(\mathbf{x}) = (V \circ \nabla v)(\mathbf{x})$ satisfy (H1), (H2) and (H3). Then, \mathbf{u}_p is continuous.*

Proof. Being both $U(\mathbf{x})$ and the objective function continuous in their arguments, Berge's theorem [3, Theorem 17.31] states that the set-valued map $\mathbf{u}_p(\mathbf{x})$ is upper hemicontinuous. However, being $U(\mathbf{x})$ convex, the projection of \mathbf{u}_i onto $U(\mathbf{x})$ is unique, meaning that \mathbf{u}_p is a singleton-valued function. Finally, for such a function, upper hemicontinuity is equivalent to continuity [3, Lemma 17.6]. \square

Corollary 4.1. *Let \mathbf{u}_i be continuous. Then, the composite control law \mathbf{u}_p , defined in (4.10) with (4.9), is continuous.*

Table 4.3 – Visual servo configurations. 3D poses are expressed as the object-to-camera pose with translation and angle-axis representation.

	Configuration	Desired pose	
		${}^o\mathbf{t}_c$	$\theta\mathbf{u}$
1	${}^oP = \begin{bmatrix} -0.5 & 0.5 & -0.5 & 0.5 \\ -0.5 & -0.5 & 0.5 & 0.5 \\ 0.0 & 0.0 & 0.0 & 0.0 \end{bmatrix}$	$[-0.59, -0.23, -0.77]$	$\left[0.0, \frac{\sqrt{2}}{2}, -\frac{\sqrt{2}}{2}\right]$
2	${}^oP = \begin{bmatrix} 0.0 & 1.0 & 0.62 & 0.31 \\ 0.0 & 0.0 & 0.38 & 0.88 \\ 0.0 & 0.0 & 0.0 & 0.31 \end{bmatrix}$	$[-0.66, -0.43, 2.49]$	$[-\pi, 0.0, 0.0]$

Proof. The proposed feasible set (4.9) trivially satisfies (H1) and, as discussed above, (H2) as long as $\epsilon < \rho$. For $\mathbf{g} = 0$, let $B_\eta(0) \subset \mathbb{R}^n$ be an open ball of radius η . By (4.8), for all $\mathbf{z} \in B_{\frac{\eta}{\rho}}(0)$ we have $V(\mathbf{z}) \subset B_\eta(0)$, meaning that (H3) is satisfied. Therefore Proposition 4.2 applies for (4.10) with (4.9). \square

4.4 Meta-controller in image-based visual servoing

In this section we illustrate some properties of the proposed controller in two different image-based visual servoing configurations. The positions of the observed points, together with the desired pose, are detailed in Table 4.3.

Scene 1 corresponds to a well-known configuration in image-based visual servoing where a local minimum is present, with 4 points forming a square [23]. Scene 2 is a case of 4 non-coplanar points, which has the same configuration of Case 1a that we saw back in Chapter 3.

All simulations are carried using ViSP software [70]. Graphs are generated with log2plot [58].

We consider in this section two of the control laws that we studied in this thesis, namely the energy-decreasing and the desired posed approximation pseudo-inverse controllers. These controllers are classically used in image-based visual servoing, and many works have been carried to combine these two approaches [23], such as using the mean between the two [67]. However, as far as the author is aware, none of these approaches are energy decreasing, making the meta-controller the first energy-decreasing combination of these well-known controllers. In our case, we rely on $\mathbf{u}_i = -\mathbf{L}_e^+ \mathbf{e}$ for its convergence properties.

This control is then projected onto a sector within the error-decrease space from (4.10).

In the simulations we compare the three approaches: current interaction matrix, desired one and meta-controller, denoted in the graphics with \mathbf{L}^+ , \mathbf{L}^{*+} and `meta`, respectively. We then compare the separation between the regions of attraction induced by \mathbf{L}_e^+ and the proposed approach. All controllers use the same λ gain for comparison purposes.

Scene 1

Using Scene 1 from Table 4.3, two initial poses are considered. The behavior from the first starting pose is depicted in Fig. 4.8 and corresponds to a classical case where using \mathbf{L}_e^+ makes the system converge to a local minimum (represented by a blue camera), while using $\mathbf{L}_{e^*}^+$ leads to the global minimum (represented by a green one). The 3D graph (Fig. 4.8a) shows that using $\mathbf{L}_{e^*}^+$ (purple trajectories) induces a non-desirable trajectory as the camera passes to the other side of the observed object. On the opposite, \mathbf{L}_e^+ (blue trajectories) leads to almost a 3D straight line, though ending in the local minimum. Using the proposed meta-controller (orange trajectories) leads to the desired pose. Figure 4.8b highlights the behavior of the proposed controller, that is energy decreasing even if the error almost stalls between 25 and 125 s. This behavior corresponds to the trajectory “sliding” along the saddle’s level set, which is the same behavior that we saw in the illustrative example above, where, as we see in Fig. 4.6c, the composite controller is able to follow an almost-constant energy level while approaching the global minimum.

Finally, Fig. 4.8c shows the difference between the meta-controller and \mathbf{u}_i , which is the result of the optimization (4.10). The control that is applied is \mathbf{u}_i for the first seconds, then moves away as \mathbf{u}_i is not energy decreasing. The largest difference is found during the stalling phase (25-100 s) when the constraints in (4.10) induce a very different behavior compared to \mathbf{u}_i . Overall, Fig. 4.8b shows that in this case the meta-controller takes more time than the desired pose approximation controller to converge to a small error level.

The same scene is now used with another starting pose in Fig. 4.9. This time all controllers converge to the desired pose. Yet, the energy-decreasing pseudo-inverse controller induces a very large 3D displacement (blue in Fig. 4.9a). Using $\mathbf{L}_{e^*}^+$ is not energy decreasing at the beginning, hence the meta-controller greatly differs from it (Fig. 4.9c). After a few seconds, the behavior induces by \mathbf{u}_i is energy decreasing and the meta-controller is the same: the objective of (4.10) can be minimized to 0.

This example highlights that the proposed meta-controller keeps the good properties of an arbitrary controller \mathbf{u}_i if the configuration is compatible. Besides, Fig. 4.9b shows that

the composite controller takes in this case less time than the desired pose approximation controller to converge to a small error level.

Scene 2


The same comparison is done on a scene with 4 non-coplanar points, as shown in Fig. 4.10. Again, the initial and final poses are chosen so that the energy-decreasing pseudo-inverse controller ends in the local minimum. This time, the meta-controller is much closer to the motion induced by \mathbf{u}_i and their trajectories are almost indistinguishable (orange and purple on Fig. 4.10a). As shown in Fig. 4.10c, the final control input is close to \mathbf{u}_i except for the last part. This slight difference is enough to make it strictly energy decreasing, while $\mathbf{L}_{e^*}^+$ is not as shown in Fig. 4.10b. This time again, the meta-controller converges faster than using the desired pose approximation controller.

Separation between the regions of attraction


In order to have a rough idea of the region of attraction induced by the meta-controller, extensive simulations have been carried out to determine the separation between the global and local regions of attraction. This is done in Scenes 1 and 2 and for the energy-decreasing pseudo-inverse controller and the meta-controller. While the separation is defined in $SE(3)$, we project it to \mathbb{R}^3 by associating to any starting camera position in \mathbb{R}^3 , the orientation that already minimizes the visual error with positive Z -depths for all the observed points. This projection reduces the separation to a manifold in \mathbb{R}^3 which is possible to display and compare. Similarly to what we did in Chapter 3, the surface that we find with this procedure is such that for any position on each side of the surface, there exists at least one corresponding orientation so that the full camera pose is inside the corresponding, true region of attraction in $SE(3)$.

The surface itself is built with the following steps:


1. Discretize the half-space where the desired pose lies;
2. For each starting point, apply the orientation that minimizes the visual error;
3. Simulate the controller and classify starting points between global or local region of attraction;
4. For each pair of points not belonging to the same region, perform a dichotomy to identify the separation;
5. Reconstruct the surface from the dichotomy results.



(a) 3D behavior



(b) Energy behavior



(c) Normalized difference between composite controller and L^{**}

Figure 4.8 – Classical local minimum configuration for pose 1 in Scene 1. The L_e^+ controller (blue) converges to a local minimum (blue) while L_e^{**} (purple) and the proposed one (orange) converge to the desired pose (green) (a). (b) log-scale of $\|e\|$, showing that the proposed controller is energy decreasing. (c) The composite controller takes a very different path compared to u_i , especially during the stalling phase (25-100 s).







Figure 4.9 – Scene 1 with another starting pose. This time all controllers converge to the desired pose (green) (a). The L_e^+ controller (blue) induces a very large 3D trajectory (blue, a) while using $L_{e^*}^+$ makes it not energy decreasing (purple, b). The proposed controller (orange) converges to the desired pose (a) while remaining energy decreasing (b). (c) The composite controller is very different from u_i at the beginning, then follows a similar direction.



(a) 3D behavior

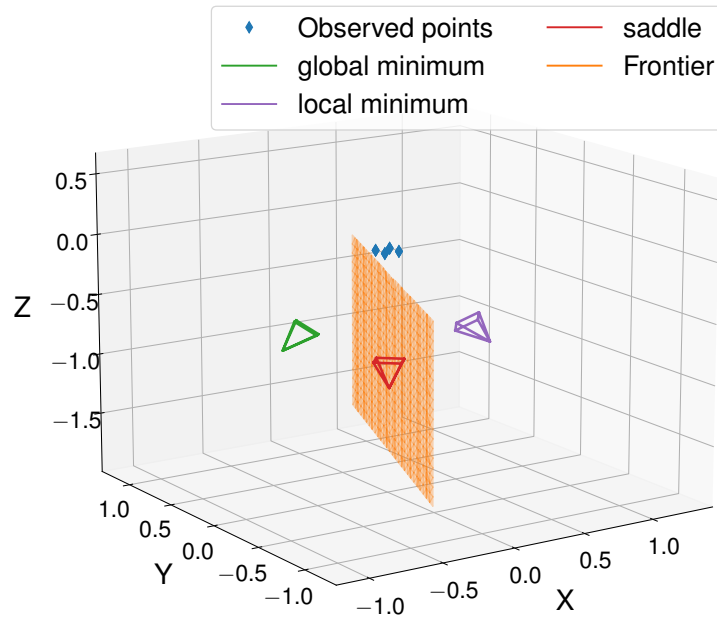


(b) Energy behavior

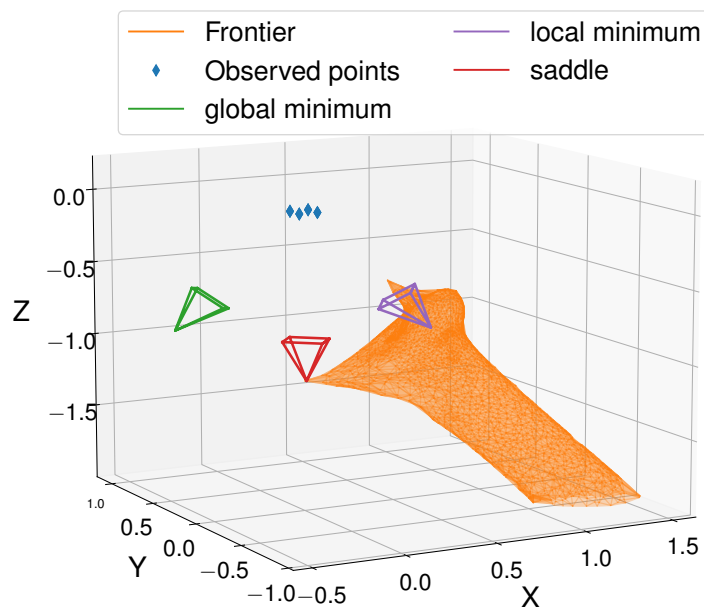


(c) Normalized difference between composite controller and L^{*+}

Figure 4.10 – Local minimum configuration in Scene 2. The L_e^+ controller (blue) converges to a local minimum while L_e^{*+} (purple) and the proposed one (orange) converge to the desired pose (a). (b) log-scale of $\|e\|$, showing that the proposed controller is energy decreasing. (c) The meta-controller is close to u_i except for the last seconds.





(a) Scene 1 with L^+




(b) Scene 1 with meta controller

Figure 4.11 – Separation between the regions of attraction of the local and global minimizers for Scene 1. Using L_e^+ (a) separates the space in two halves, each converging to a given minimizer. Using the proposed meta-controller (b) the region of attraction of the local minimizer is much smaller. In each case the saddle point (red) lies on the frontier.


(a) Scene 2 with L^+ 

(b) Scene 2 with meta controller

Figure 4.12 – Separation between the regions of attraction of the local and global minimizers for Scene 2. Using L_e^+ (a) separates the space in two halves, each converging to a given minimizer. Using the proposed meta-controller (b) the region of attraction of the local minimizer is much smaller. In each case the saddle point (red) lies on the frontier.



(a) Scene 1



(b) Scene 2

Figure 4.13 – Comparison of the separation between the regions of attraction and the component connected to the local minimum of the energy’s level set, computed at the saddle’s energy level, for Scene 1 (a) and Scene 2 (b).

Figures 4.11-4.12 shows the resulting separation. The classical local minimum configuration is displayed in Fig. 4.11a where the separation for \mathbf{L}_e^+ is a plane, meaning that each minimizer has roughly the same region of attraction. This is also the case for non-coplanar points in Scene 2 as shown in Fig. 4.12a, even if the separation is a deformed plane.

On the opposite, using the composite controller reduces the separation to a much smaller surface, as shown in Fig. 4.11b for Scene 1 and Fig. 4.12b for Scene 2. Similarly to what we had in the meta-controller’s illustrative example (Section 4.3.1), the region of attraction of the local minimizer is the union of two volumes:

- The closed region around the local minimizer that has an energy smaller than the saddle point (i.e., the component of the sublevel set which is connected to the local minimizer);
- and an infinite branch corresponding to starting positions with an error greater than the saddle point, that still converge to the local minimizer.

We can see a comparison between the separating surfaces and the sublevel set’s component in Figure 4.13. The similarity with Figure 4.6c (again in Section 4.3.1) is striking: we can see that, in a close proximity of the saddle point, its stable manifold is non-smooth and it “envelopes” the level set’s component, meaning that trajectories starting close to this boundary will slide along its surface. We want to highlight that the procedures we employed to compute the level sets and the separating surfaces are very different, so we find remarkable the almost perfect overlap between the two, which substantiates our findings.

In conclusion, while the level set’s component around the local minimum represents an inescapable region for any energy-decreasing controller, the infinite branch which completes its region of attraction depends on the employed control strategy, and we show that it is much smaller for the meta-controller than for the classical energy-decreasing pseudo-inverse approach.

HANDLING UNCERTAINTIES

In this last chapter, we will provide a preliminary investigation of the impact of uncertainties on the equilibria of energy-decreasing dynamical systems.

Typically, uncertainties in image-based visual servoing come mainly from two different sources, either from an imprecise knowledge of the camera intrinsic parameters or from coarse estimation of the points' depths [22]. The impact of these errors on the equilibria and their stability is not well understood. Most works on the topic focus on the stability of the global equilibrium alone, using simplified models of the camera dynamics [24, 37, 38]. In [69], it is shown that, again for the global equilibrium alone, a relatively low estimation error on the points' depths can make the desired pose unstable. Being a formal study of the stability of these systems very complicated, it seems that later works in the field took a different direction and focused mainly on the choice of different visual features.

No results that we are aware of treat the case of local minima, and it is quite easy to guess why this is the case: even if we assume to have a perfect knowledge of the camera intrinsic parameters (and, thus, the points' Cartesian coordinates), we cannot expect the equilibria to be robust with respect to uncertainties, since the interaction matrix depends on the points' depths as well. An uncertainty on these depths will in particular change the kernel of \mathbf{L}_e^\top , meaning that, in general, this uncertainty will impact all the solutions of the equilibrium condition $\mathbf{L}_e^\top \mathbf{e} = 0$, except the trivial one $\mathbf{e} = 0$. An imprecise estimation of the camera intrinsic parameters additionally worsen the situation, making the whole concept of equilibrium quite vacuous.

In an attempt to make sense of how to deal with uncertainties in image-based visual servoing and energy-decreasing systems at large, we started to explore two strategies. The first one is to characterize quantitatively what is the impact of unmodeled perturbations on the closed-loop system. Intuitively, it is reasonable to expect that, if the perturbations are bounded and small enough, the system will still try to converge towards a stable equilibrium. Once it arrives in a neighborhood of the equilibrium, though, we also expect the vector field to become negligible with respect to the perturbation, meaning that the

system stays in the neighborhood but it is actually incapable of converging to the stable equilibrium¹. In this settings, we say that the system is *practically stable*, and we propose a strategy to find explicitly the *invariant* neighborhood of the stable equilibrium towards which we are guaranteed to converge.

The second and final idea that we explored to deal with uncertainties in the closed-loop system is to employ a control strategy which is *robust* to perturbation, i.e., such that the stable equilibria do not change under (reasonable) errors in the model's parameters. The control strategy that we considered is the so-called *super-twisting algorithm*, which is a sliding mode controller with well-known robustness and finite-time convergence properties. Our contributions on this last approach are limited and very technical (we extended known convergence conditions, allowing for an uncertain input gain), and, perhaps, they are the less applicable to the visual servoing context at the current state.

Since we were not able to flesh out the impact of uncertainties in image-based visual servoing, we focus in this chapter on the methodological aspects of handling uncertainties, testing out our approaches on simple examples. Nevertheless, we believe that the techniques that we studied, and the practical stability assessment in particular, could be effective in image-based visual servoing as well.

The chapter is organized as follows. In Section 5.1, we introduce the practical stability of perturbed dynamical system and we define the *value function* associated to the system, which, as we will see, is the solution of a parametric optimization problem which allows us to explicitly compute an invariant approximation of the system's attractors and their region of attraction. We then study the value function's continuity, which will prove itself to be a crucial property to simplify its computation, and we apply this method to a class of perturbed linear system. Then, in Section 5.2, we introduce the super-twisting algorithm and we show that, under a suitable change of variable, it is possible to guarantee finite-time convergence to the origin even when the control input is multiplied by an unknown, time-varying gain.

1. This intuition fails if the stable equilibrium is *robust* to perturbations. For instance, this is the case, in the visual servoing context, for the desired pose when only the points' depths are perturbed [69]. Nevertheless, as we will see, our methodology is still applicable in these cases, and it simply provides a single point (the equilibrium) as the attractor's estimate.

5.1 Practical stability of perturbed dynamical systems

The estimation of regions of attraction for perturbed dynamical systems is an old problem in the literature, and methods based on Lyapunov theory are probably the most wide-spread [48, 55, 60]. As we saw back in Chapter 1, the main idea behind these approaches is to estimate the region of attraction as the invariant sub-level sets of a given Lyapunov function. However, all these methods require that there exists a stable and robust (i.e. perturbation-independent) equilibrium for the system. This is a major drawback, since in a large class of real-world problems this hypothesis is not satisfied. An important example of this is the problem of tracking a moving target with unknown dynamics, which is a common problem that cannot, a priori, allow for a robust equilibrium point. In these problems it could still be possible however to discuss the *practical stability* of the system, i.e. to assess whether the system, while technically unstable in theory, can be considered to be stable in practice. Consider, for instance, the case where the tracking system quickly reaches a small neighborhood of the target while never actually converging to it. Motivated by this, we want to relax the requirement of a robust equilibrium by using instead the concept of *attractor*. An attractor is a set of states which is positively invariant and attracting, i.e. such that the system converges to it (possibly in infinite time) when its initial state is close enough to the attractor. In the tracking system example discussed above, the attractor would be the small neighborhood of the target; this characterization, however, holds in more general settings.

The main contribution of this section is to characterize the practical stability of perturbed systems with bounded perturbations by exploiting a given candidate Lyapunov function $v(\mathbf{x})$. The proposed definition of *value function* associated to the system and candidate Lyapunov function allows characterizing the system's attractors, hence extending the classical Lyapunov analysis of attracting fixed points. Our formulation does not require any underlying structure for the dynamics of the system, relaxing the classical hypothesis of polynomial dynamics and allowing for the study of a large class of non-linear, perturbed systems.

In practice, one can think of v to be the *energy function* for the nominal, unperturbed system. In this case, we can then see our method as a characterization of the perturbation's effects on the nominal controller, which we could use in image-based visual servoing as well. Nevertheless, this is not a strict requirement for our analysis, and we will not restrict

to these settings throughout the section.

The section is organized as follows. In Section 5.1.1 we generalize the classical region of attraction’s estimation problem to the identification of an attractor and the estimation of its region of attraction, and we show that this problem can be framed as a parametric optimization problem, using the candidate Lyapunov function’s level sets as parameter. In Section 5.1.2 we study the continuity of the parametric optimization problem’s solution, showing that, under common assumptions, we can expect it to be well-behaved. Finally, in Section 5.1.3 we apply our results to discuss the practical stability of a class of perturbed linear systems.

5.1.1 Value function definition

We consider a perturbed system $\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{w})$, where $\mathbf{w} \in \mathcal{W}$ represents a time-dependent, unknown perturbation, with f locally Lipschitz with respect to \mathbf{x} and continuous with respect to \mathbf{w} . We further assume that \mathcal{W} is a compact set, and we define a *candidate Lyapunov function*, that we denote with $v(\mathbf{x})$, to be a continuously differentiable function satisfying:

$$v(\mathbf{x}) \geq 0 \quad \text{and} \quad v(\mathbf{x}_0) = 0 \quad (5.1)$$

for some \mathbf{x}_0 .

Given f and v , we propose to define the function $m(c)$ that associates to a target value c of the candidate Lyapunov function the worst-case Lie derivative on the corresponding level set:

$$m(c) = \sup_{\mathbf{x} \in \mathcal{V}_c} \Phi(\mathbf{x}) \quad (5.2a)$$

$$\Phi(\mathbf{x}) = \sup_{\mathbf{w} \in \mathcal{W}} \nabla v(\mathbf{x})^\top f(\mathbf{x}, \mathbf{w}), \quad (5.2b)$$

where $\mathcal{V}_c = \{\mathbf{x} \in \mathcal{X} : v(\mathbf{x}) = c\}$ is the *level set*² with value c . We call this function the *value function*³ associated to f and v , or simply the value function when the associated system and candidate Lyapunov function are clear from the context. Sufficient conditions for the two supremum to be maximum are investigated in the next section. The obvious

2. Please notice that this notation is different from the one introduced in Chapter 1, where \mathcal{V}_c denotes the *sublevel sets*.

3. Most of the theoretical machinery used throughout the section comes mainly from economics, the field in which the name *value function* has been defined.

usefulness of the value function is summarized in the following two properties, valid under the typical assumption that the level sets of the candidate Lyapunov function involved in these properties are compact:

1. If $m(c) < 0$ inside the interval $(\underline{c}, \bar{c}]$, for some $0 < \underline{c} < \bar{c}$, then the sublevel set $\mathcal{V}_{\leq \underline{c}}$ is an attractor and the sublevel set $\mathcal{V}_{\leq \bar{c}}$ is inside its region of attraction, i.e. $\mathcal{V}_{\leq \bar{c}}$ is an *estimated region of attraction* for the attractor.
2. If $m(c) < 0$ inside $(\underline{c}, +\infty)$ for some $0 < \underline{c}$, then the sublevel set $\mathcal{V}_{\underline{c}}$ is a global attractor.

As discussed in the introduction, for a *practically stable* system we typically expect that $m(c) > 0$ inside $(0, c^*)$ and $m(c) < 0$ inside $(c^*, +\infty)$: provided that v has bounded level sets, all trajectories will converge inside the positively invariant set \mathcal{V}_{c^*} . The usefulness of the value function defined here is of course balanced by the difficulty of computing it. Its formal evaluation is carried out for a class of perturbed linear system in Section 5.1.3.

In the general case, the quest of finding an interval where the value function is negative provides a strong temptation: computing its roots is of course a simpler and appealing approach. The roots of the value function can be characterized by using the Karush–Kuhn–Tucker (KKT) conditions applied to the optimization problem⁴ (5.2):

$$\nabla^2 v(\mathbf{x}) f(\mathbf{x}, \mathbf{w}) + (f_{\mathbf{x}}(\mathbf{x}, \mathbf{w})^\top + \lambda \mathbf{I}) \nabla v(\mathbf{x}) = 0 \quad (5.3a)$$

$$(f_{\mathbf{w}}(\mathbf{x}, \mathbf{w})^\top + \mu \mathbf{I}) \nabla g(\mathbf{w}) = 0 \quad (5.3b)$$

$$v(\mathbf{x}) - c = 0 \quad (5.3c)$$

$$\mu g(\mathbf{w}) = 0 \quad (5.3d)$$

$$\nabla v(\mathbf{x})^\top f(\mathbf{x}, \mathbf{w}) = 0, \quad (5.3e)$$

where $\nabla^2 v$ denotes the Hessian of v , $f_{\mathbf{x}}$, $f_{\mathbf{w}}$ the Jacobians of f w.r.t. \mathbf{x} and \mathbf{w} , $g(\mathbf{w})$ represents the constraints on \mathbf{w} (i.e. $\mathcal{W} = \{\mathbf{w} : g(\mathbf{w}) \leq 0\}$) and $\lambda, \mu \in \mathbb{R}$ are the KKT multipliers. In addition to the usual KKT conditions, we consider c as a variable and we add the constraint $\dot{v}(\mathbf{x}, \mathbf{w}) = 0$ in order to find a solution compatible with $m(c) = 0$. However, this simple intuition hides a trap: if the value function is not continuous, we might incur in a change of sign without passing through zero, rendering this approach useless. Thus, before we succumb to temptation, it is crucial to characterize the properties

4. By substituting the definition of $\Phi(\mathbf{x})$ into (5.2a), we can see $m(c)$ as the solution of a single optimization problem in \mathbf{x} and \mathbf{w} .




Figure 5.1 – Example of upper (on the left) and lower (on the right) semicontinuous functions, where the solid black point denotes $h(1)$. In practice, the only difference is whether $h(1)$ is at the top or bottom portion of the discontinuity.

of f, v that guarantee the value function’s continuity.

5.1.2 Continuity of the value function

As seen in the previous section, the continuity of the value function plays an important role in choosing the strategy to evaluate it. We want to study the continuity of $m(c)$ by investigating its *upper semicontinuity* and *lower semicontinuity* separately. Intuitively, the continuity property can be “split” into these two sub-properties, and we find that a function is continuous when both hold at the same time. For a scalar function (like $m(c)$), the following definition applies [3].

Definition 5.1. Let $h : \mathcal{H} \subset \mathbb{R} \rightarrow \mathbb{R} \cup \{-\infty, +\infty\}$. Then:

1. The function h is called *upper semicontinuous* at \bar{c} if, for any $\epsilon > 0$, there exists $\delta > 0$ such that:

$$|c - \bar{c}| < \delta \implies h(c) \leq h(\bar{c}) + \epsilon \tag{5.4}$$

2. The function h is called *lower semicontinuous* at \bar{c} if, for any $\epsilon > 0$, there exists $\delta > 0$ such that:

$$|c - \bar{c}| < \delta \implies h(c) \geq h(\bar{c}) - \epsilon \tag{5.5}$$

A function h which is upper (lower) semicontinuous on its whole domain is simply called *upper (lower) semicontinuous*. If it is both upper and lower semicontinuous, then it is a *continuous function*.

An example of the difference between upper and lower semicontinuity can be seen in Figure 5.1. While, technically, in this definition we ask the codomain of h to be the

extended reals $\mathbb{R} \cup \{-\infty, +\infty\}$, in practice, for $m(c)$, we can consider to have the set of real numbers \mathbb{R} as codomain. Moreover, due to the positive definiteness of v , its domain is \mathbb{R}_+ .

We are ready to discuss our results. First of all, we show that, under typical assumptions on f and v , the function $\Phi(\mathbf{x})$ is continuous.

Proposition 5.1. *Let $v : \mathcal{X} \rightarrow \mathbb{R}$ be a function of class \mathcal{C}^1 , and let $f : \mathcal{X} \times \mathcal{W} \rightarrow \mathbb{R}^n$ be continuous in \mathbf{x} and \mathbf{w} , where \mathcal{W} is a compact set. Then, $\Phi(\mathbf{x})$ is a continuous function and the maximum is attained, i.e. $\Phi(\mathbf{x})$ can be written as a max instead of sup.*

Proof. Being both $\nabla v(\mathbf{x})$ and $f(\mathbf{x}, \mathbf{w})$ continuous, the product $\nabla^\top v(\mathbf{x})f(\mathbf{x}, \mathbf{w})$ is continuous as well. Thus, by considering \mathcal{W} as a constant (and so continuous) set-valued function with nonempty compact values, we can directly apply Berge's maximum theorem [3, Theorem 17.31], which proves that $\Phi(\mathbf{x})$ is continuous. \square

While promising, the continuity of $\Phi(\mathbf{x})$ is not sufficient to guarantee the continuity of the value function. Consider, for instance, the unperturbed, scalar system defined as:

$$\dot{x} = f(x) = \begin{cases} -\frac{10x}{(0.1x^2+5)^2} & \text{if } x > 0 \\ -2x & \text{if } x \leq 0 \end{cases}, \quad (5.6)$$

and let:

$$v(x) = \begin{cases} \frac{x^2}{0.1x^2+5} & \text{if } x > 0 \\ x^2 & \text{if } x \leq 0 \end{cases} \quad (5.7)$$

be its candidate Lyapunov function. The plot of v is shown in Figure 5.2.

It is easy to verify the continuity and continuous differentiability of f and v , respectively. Moreover, we can see that $f(x) = -v'(x)$, meaning that the Lie derivative is $\dot{v}(x) = -(v'(x))^2$, which is negative for all $x \in \mathbb{R} \setminus \{0\}$. Even with a scalar and globally asymptotically stable system, we can see, in Figure 5.2, that the value function is not upper semicontinuous: this is due to a sudden change of the worst-case Lie derivative at the level of the horizontal asymptote.

The lower semicontinuity of the value function is not guaranteed either. Consider, for instance, the system:

$$\dot{x} = f(x) = -4x(x^2 - x - 2), \quad (5.8)$$




Figure 5.2 – On the left, plot of $v(x)$ defined in (5.7). The dotted line represents the horizontal asymptote for $x \rightarrow +\infty$. On the right, the corresponding value function, where the solid black point denotes the value of m at the discontinuity.




Figure 5.3 – On the left, plot of $v(x)$ defined in (5.9). On the right, the corresponding value function, where the solid black points denote the values of m at the discontinuities.

with candidate Lyapunov function:

$$v(x) = x^4 - \frac{4}{3}x^3 - 4x^2 + \frac{32}{3}. \quad (5.9)$$

Similarly to the previous example, we have a system of the form $f(x) = -v'(x)$. As we can see in Figure 5.3, in this case as well, the value function is discontinuous, in particular it is not lower semicontinuous: this is due to sudden changes of the worst-case Lie derivative at the levels of the local minimizer and local maximizer.

Motivated by this, we want, in the following theorem, to independently study the upper and lower semicontinuity of the value function. The theorem’s proof relies on the same foundations of the already introduced Berge’s maximum theorem: as we will see, the upper and lower semicontinuity of $m(c)$ depend on the continuity of $\Phi(\mathbf{x})$ and the

hemicontinuity⁵ of \mathcal{V}_c , seen as a set-valued function $\mathcal{V}_c : \mathbb{R}_+ \rightrightarrows \mathcal{X}$. We have the following.

Theorem 5.1. *Let v, f satisfy the same hypotheses of Proposition 5.1. Then, the following holds true:*

1. *If v is radially unbounded, then $m(c)$ is upper semicontinuous and the maximum is attained, i.e. $m(c)$ can be written as a max instead of sup.*
2. *If, for all $\mathbf{x} \in \mathcal{V}_{\bar{c}}$, $\nabla v(\mathbf{x}) \neq 0$, then $m(c)$ is lower semicontinuous at \bar{c} .*

Remark. *Case 1 rules out situations illustrated by (5.6)–(5.7), while Case 2 rules out situations illustrated by (5.8)–(5.9).*

Proof. Due to Proposition 5.1, $\Phi(\mathbf{x})$ is continuous in the theorem’s hypotheses. Thus, we focus in this proof on the upper and lower hemicontinuity of \mathcal{V}_c .

We start with the first statement. Being v continuous, then the preimage of a closed set through v is closed as well, meaning that \mathcal{V}_c is closed. Moreover, the radial unboundness of v implies that \mathcal{V}_c is bounded, for all $c \in \mathbb{R}_+$. Thus, by [11, Corollary 21], \mathcal{V}_c is upper hemicontinuous, which means that, by [3, Lemma 17.30], $m(c)$ is a max and it is upper semicontinuous.

The second statement’s proof is proved in a similar fashion, albeit more involved. Firstly, we show that \mathcal{V}_c is *inner hemicontinuous* at \bar{c} , which is a property equivalent to the lower hemicontinuity [11, Proposition 23]. By its definition, we have that \mathcal{V}_c is inner hemicontinuous at \bar{c} if, for all $\bar{\mathbf{x}} \in \mathcal{V}_{\bar{c}}$, it holds true that:

$$\forall \{c_k\} \rightarrow \bar{c}, \exists \mathbf{x}_k \in \mathcal{V}_{c_k} : \mathbf{x}_k \rightarrow \bar{\mathbf{x}} \quad (5.10)$$

We prove the inner hemicontinuity of \mathcal{V}_c by contradiction. Suppose that \mathcal{V}_c is not inner hemicontinuous at \bar{c} . That means that there exists $\bar{\mathbf{x}} \in \mathcal{V}_{\bar{c}}$ such that:

$$\exists \{c_k\} \rightarrow \bar{c} : \forall \mathbf{x}_k \in \mathcal{V}_{c_k}, \mathbf{x}_k \not\rightarrow \bar{\mathbf{x}}. \quad (5.11)$$

Consider the restriction of v , centered in $\bar{\mathbf{x}}$ and along the direction of its gradient evaluated at $\bar{\mathbf{x}}$, i.e. the scalar function $\bar{v}(h) = v(\bar{\mathbf{x}} + h\nabla v(\bar{\mathbf{x}}))$. Notice that, by the continuity of v , \bar{v} is continuous. Moreover, there exists an open interval $H \subset \mathbb{R}$, with $0 \in H$, such that \bar{v} is strictly increasing over H . This is due to the fact that the derivative $D\bar{v}(h)$ evaluated at zero is:

$$D\bar{v}(h)|_{h=0} = \|\nabla v(\bar{\mathbf{x}})\|^2 > 0, \quad (5.12)$$

5. See Definition 4.1 in Section 4.3.3.

and, by the continuity of ∇v , there exists a neighbourhood of the origin (i.e. an open interval H containing zero) where (5.12) holds true, meaning that \bar{v} is strictly increasing on H . Thus, \bar{v} admits a continuous inverse \bar{v}^{-1} on H . Its continuity means that:

$$\forall \{c_k\} \rightarrow \bar{c}, \quad \bar{v}^{-1}(c_k) \rightarrow \bar{v}^{-1}(\bar{c}) = 0 \quad (5.13)$$

Let $h_k = \bar{v}^{-1}(c_k)$. By defining the sequence $\{\mathbf{x}_k\}_{k \in \mathbb{N}}$ as $\mathbf{x}_k = \bar{\mathbf{x}} + h_k \nabla v(\bar{\mathbf{x}})$, it follows from (5.13) that:

$$\forall \{c_k\} \rightarrow \bar{c}, \quad \mathbf{x}_k \rightarrow \bar{\mathbf{x}} \quad (5.14)$$

which is in contradiction with 5.11, proving the inner (and, thus, lower) hemicontinuity of \mathcal{V}_c at \bar{c} . By [3, Lemma 17.29], the lower hemicontinuity of \mathcal{V}_c at \bar{c} implies that $m(c)$ is lower semicontinuous at \bar{c} , concluding the proof. \square

5.1.3 Application to linear tracking systems with bounded velocity target

Upper bound of the value function for linear systems with bounded uncertain right hand side

We consider the linear system


$$\dot{\mathbf{z}} = \mathbf{A}\mathbf{z} + \mathbf{w}, \quad (5.15)$$

where \mathbf{A} is a *stable* matrix (i.e. all eigenvalues have negative real part) and the uncertainty \mathbf{w} is bounded inside $\mathcal{W} = \{\mathbf{w} \in \mathbb{R}^n : \|\mathbf{w}\| \leq M\}$. We consider a quadratic Lyapunov function $v(\mathbf{z}) = \mathbf{z}^\top \mathbf{P}\mathbf{z}$ for the nominal system with \mathbf{P} symmetric positive-definite (SPD). The Lie derivative for the system with nominal $\mathbf{w} = 0$ is given by $\mathbf{z}^\top (\mathbf{P}\mathbf{A} + \mathbf{A}^\top \mathbf{P})\mathbf{z}$, which is supposed to be negative-definite. Such a matrix \mathbf{P} is usually obtained by choosing an arbitrary SPD matrix \mathbf{Q} , typically $\mathbf{Q} = \mathbf{I}$, and solving the Lyapunov equation

$$\mathbf{P}\mathbf{A} + \mathbf{A}^\top \mathbf{P} + \mathbf{Q} = 0. \quad (5.16)$$

The following theorem provides an upper bound for the value function associated to the system and the Lyapunov function defined above.

Theorem 5.2. *Let \mathbf{P} , \mathbf{Q} be positive-definite matrices so that Lyapunov's equation (5.16)*


 Figure 5.4 – Graph of $-c + \sqrt{c}$.

holds. Then, $m(c) \leq \bar{m}(c)$ with

$$\bar{m}(c) = \sigma_1 c + 2M\sqrt{\lambda_1}\sqrt{c}, \quad (5.17)$$

where σ_1 is the greatest eigenvalue of $(-\mathbf{P}^{-1}\mathbf{Q})$ and λ_1 is the greatest eigenvalue of \mathbf{P} . Furthermore, $\sigma_1 < 0$.

Proof. With these hypotheses, $m(c)$ can be written as:

$$m(c) = \max_{\mathbf{z} \in \mathcal{V}_c} \Phi(\mathbf{z}). \quad (5.18)$$

Here we have

$$\Phi(\mathbf{z}) = \max_{\mathbf{w} \in \mathcal{W}} (-\mathbf{z}^\top \mathbf{Q}\mathbf{z} + 2\mathbf{z}^\top \mathbf{P}\mathbf{w}) \quad (5.19)$$

$$= -\mathbf{z}^\top \mathbf{Q}\mathbf{z} + 2 \max_{\mathbf{w} \in \mathcal{W}} \mathbf{z}^\top \mathbf{P}\mathbf{w}. \quad (5.20)$$

Using Cauchy-Schwarz inequality and the bound on the norm of \mathbf{w} , we obtain $\mathbf{z}^\top \mathbf{P}\mathbf{w} \leq \|\mathbf{P}\mathbf{z}\| \|\mathbf{w}\| \leq M\|\mathbf{P}\mathbf{z}\|$, with the equality holding when \mathbf{w} is parallel to $\mathbf{P}\mathbf{z}$. Since \mathcal{W} is

a ball of radius M , such a \mathbf{w} parallel to \mathbf{Pz} with norm M exists and we have $\Phi(\mathbf{z}) = -\mathbf{z}^\top \mathbf{Qz} + 2M\|\mathbf{Pz}\|$. This leads us to:

$$m(c) = \max_{\mathbf{z} \in \mathcal{V}_c} \left[-\mathbf{z}^\top \mathbf{Qz} + 2M\|\mathbf{Pz}\| \right] \quad (5.21a)$$

$$\leq \left(\max_{\mathbf{z} \in \mathcal{V}_c} -\mathbf{z}^\top \mathbf{Qz} \right) + 2M \left(\max_{\mathbf{z} \in \mathcal{V}_c} \|\mathbf{Pz}\| \right), \quad (5.21b)$$

where the optimization problem is split into two independent subproblems that can be solved independently.

For the second subproblem, we solve $\max_{\mathbf{z} \in \mathcal{V}_c} \|\mathbf{Pz}\|^2$. Lagrange's first-order conditions yield $\mathbf{P}^2 \mathbf{z} - \mu \mathbf{Pz} = 0$, which must hold for some $\mu \in \mathbb{R}$. Being \mathbf{P} invertible, this is equivalent to solving $\mathbf{Pz} = \mu \mathbf{z}$, leading to

$$\max_{\mathbf{z}^\top \mathbf{Pz} = c} \|\mathbf{Pz}\|^2 = \max_{\substack{\mathbf{z}^\top \mathbf{Pz} = c \\ \mu \in \Lambda(\mathbf{P})}} \mu \mathbf{z}^\top \mathbf{Pz} = \lambda_1 c, \quad (5.22)$$

where $\Lambda(\mathbf{P})$ is the spectrum of \mathbf{P} and $\lambda_1 = \max \{\Lambda(\mathbf{P})\}$ is the greatest eigenvalue of \mathbf{P} . The second subproblem maximum is therefore $\sqrt{\lambda_1 c}$.

For the first subproblem, Lagrange's first-order conditions of $\max_{\mathbf{z} \in \mathcal{V}_c} -\mathbf{z}^\top \mathbf{Qz}$ lead to $-\mathbf{Qz} - \mu \mathbf{Pz} = 0$ for some $\mu \in \mathbb{R}$. We notice that the values of μ that satisfy this equation are the generalized eigenvalues [83] of the symmetric matrix pencil $(-\mathbf{Q}, \mathbf{P})$, whose set we denote with $\Lambda(-\mathbf{Q}, \mathbf{P})$. Being \mathbf{P} positive-definite, we know that we have n real generalized eigenvalues (see [83, Theorem 15.3.3]), that we denote, w.l.o.g., with $\sigma_1 \geq \dots \geq \sigma_n$. Thus, by substitution, we have that:

$$\max_{\mathbf{z}^\top \mathbf{Pz} = c} -\mathbf{z}^\top \mathbf{Qz} = \max_{\substack{\mathbf{z}^\top \mathbf{Pz} = c \\ \mu \in \Lambda(-\mathbf{Q}, \mathbf{P})}} \mu \mathbf{z}^\top \mathbf{Pz} = \sigma_1 c, \quad (5.23)$$

which allows us to write:

$$m(c) \leq \sigma_1 c + 2M\sqrt{\lambda_1 c}. \quad (5.24)$$

We conclude the proof by noticing that, being \mathbf{P} invertible, the generalized eigenvalues of the pencil $(-\mathbf{Q}, \mathbf{P})$ are equal to the eigenvalues of $-\mathbf{P}^{-1}\mathbf{Q}$, and that, being $-\mathbf{z}^\top \mathbf{Qz} < 0$ for all $\mathbf{z} \neq 0$, we have that $\sigma_1 c < 0$ for $c \neq 0$, proving that σ_1 is indeed negative. \square

If M is not zero then $\bar{m}(c)$ starts with an infinite derivate at $c = 0$, has a single maximum at $c = \lambda_1 \left(\frac{M}{\sigma_1}\right)^2$ and $\lim_{c \rightarrow \infty} \bar{m}(c) = -\infty$. Its typical graph is shown in Figure 5.4.

The function $\bar{m}(c)$ has a unique positive root

$$c^* = \frac{4M^2\lambda_1}{\sigma_1^2}, \quad (5.25)$$

and is negative for greater values of c . The positive definite quadratic Lyapunov function being radially unbounded, the sublevel set $\mathcal{V}_{\leq c^*}$ is proved to be a global attractor of the uncertain system (5.15). In the typical case where $\mathbf{Q} = \mathbf{I}$ we have $\sigma_1 = \frac{-1}{\lambda_1}$ and the root's expression simplifies to

$$c^* = 4M^2\lambda_1^3. \quad (5.26)$$

Linear tracking systems with bounded velocity target

We consider a simple linear system with known input matrix gain \mathbf{A} of the form:

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{u}, \quad (5.27)$$

where $\mathbf{x} \in \mathbb{R}^n$ is the system's state, $\mathbf{u} \in \mathbb{R}^n$ its input and $\mathbf{A} \in \mathbb{R}^{n \times n}$ is stable. The target reference signal $\mathbf{r}(t) \in \mathbb{R}^n$ is measured at all times, and we assume the knowledge of an upper bound M on $\|\dot{\mathbf{r}}(t)\|$. The proportional control $\mathbf{u}(t) = -k(\mathbf{x}(t) - \mathbf{r}(t))$, $k > 0$, leads to the following closed loop system:

$$\dot{\mathbf{x}}(t) = k\mathbf{A}(\mathbf{x}(t) - \mathbf{r}(t)). \quad (5.28)$$

The dynamic of the tracking error $\mathbf{z}(t) = \mathbf{x}(t) - \mathbf{r}(t)$ is then

$$\dot{\mathbf{z}}(t) = k\mathbf{A}\mathbf{z}(t) - \dot{\mathbf{r}}(t). \quad (5.29)$$

We use Theorem 5.2 to find a globally attracting sublevel set $\mathcal{V}_{\leq c^*}$. In order to compare the sublevel sets for different gains, we consider the Lyapunov function $V(\mathbf{z}) = \mathbf{z}^\top \mathbf{P}\mathbf{z}$ with $\mathbf{P}\mathbf{A} + \mathbf{A}^\top \mathbf{P} + \mathbf{I} = 0$ independently of the gain k . Therefore, the matrix \mathbf{Q} that satisfies the Lyapunov equation (5.16) is $\mathbf{Q} = k\mathbf{I}$. Since \mathbf{Q} is proportional to \mathbf{I} , the sublevel value (5.25) now simplifies to

$$c^* = \frac{4M^2\lambda_1^3}{k^2}. \quad (5.30)$$




Figure 5.5 – Left: $\mathbf{r}(t)$; right: $\mathbf{r}'(t)$ with a circle of radius 6.

Table 5.1 – Values of c^* obtained using Theorem 5.2 for different gain values.

k	5	10	50	100
c^*	3.73	0.932	0.0373	0.00931

Numerical application

We consider the following highly oscillating target,

$$\mathbf{r}(t) = \begin{pmatrix} \cos(2t) + 2 \sin(2t) + 0.5 \sin(3t) \\ \sin(t) - \cos(3t) \end{pmatrix}, \quad (5.31)$$

whose velocity is bounded by $M = 6$, see Figure 5.5. The matrix \mathbf{A} is chosen to be non symmetric and \mathbf{P} follows solving $\mathbf{P}\mathbf{A} + \mathbf{A}^\top \mathbf{P} + \mathbf{I} = 0$:

$$\mathbf{A} = \begin{pmatrix} -2 & 5 \\ -1 & -1 \end{pmatrix} \quad \text{and} \quad \mathbf{P} = \begin{pmatrix} \frac{3}{14} & \frac{1}{14} \\ \frac{1}{14} & \frac{6}{7} \end{pmatrix}. \quad (5.32)$$

The largest eigenvalue of \mathbf{P} is $\lambda_1 = \frac{1}{28}(15 + \sqrt{85}) \approx 0.86$. The globally attractive sublevel values obtained using Theorem 5.2 for different values of the gain k are given in Table 5.1. The trajectories of the system and the globally attracting sublevel sets obtained by Theorem 5.2 are shown in Figure 5.6, while in Figure 5.7 we can see the value of the Lyapunov function with respect to time.

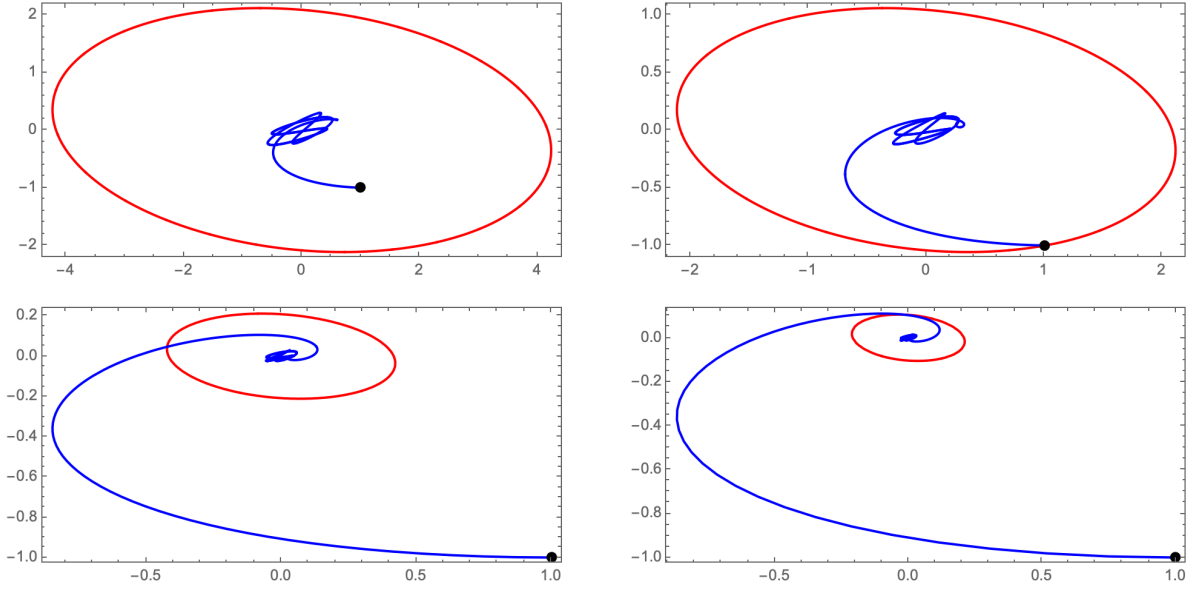


Figure 5.6 – Tracking error of the proportional linear tracking system for the same initial condition, whose initial error is denoted with a black point, together with the globally attracting ellipsoids, for different values of the gain $k \in \{5, 10, 50, 100\}$, with western reading direction.

Application of Theorem 5.1

The positive definite quadratic Lyapunov function satisfies the hypothesis of Theorem 5.1, which proves that the exact value function $m(c)$ is continuous. We can therefore infer the sign of $m(c)$ by computing its roots. The system (5.3) that characterizes its roots is now

$$-\mathbf{Q}\mathbf{x} + \mathbf{P}\mathbf{w} + \lambda\mathbf{P}\mathbf{x} = 0 \quad (5.33a)$$

$$\mathbf{P}\mathbf{x} + \mu\mathbf{w} = 0 \quad (5.33b)$$

$$\mathbf{x}^\top \mathbf{P}\mathbf{x} = c \quad (5.33c)$$

$$\mu(\mathbf{w}^\top \mathbf{w} - M^2) = 0 \quad (5.33d)$$

$$-\mathbf{x}^\top \mathbf{Q}\mathbf{x} + 2\mathbf{x}^\top \mathbf{P}\mathbf{w} = 0. \quad (5.33e)$$

It is polynomial of degree three⁶ and can be solved using formal computations. For each gain value, we find four solutions showing two distinct values of c , given in Table 5.2. The

6. The only degree three monomial in the system is the complementarity constraint (5.33d), which is actually a simple alternative. The system can therefore be solved as two degree two systems.




Figure 5.7 – Lyapunov function's value along the trajectories of the proportional linear tracking system (in blue) together with the globally attracting level set's value (in orange), for different values of the gain $k \in \{5, 10, 50, 100\}$, with western reading direction.

Table 5.2 – Values of c for the extremal values leading to a zero Lie derivative, with the corresponding values of the value function.

k	5	10	50	100
c_1	0.0507	0.0127	0.000507	0.000127
c_2	3.73	0.932	0.0373	0.00931
$m(c_1)$	2.21	1.10	0.221	0.110
$m(c_2)$	0	0	0	0

solutions with c_2 correspond exactly to the solutions previously computed, while $c_1 < c_2$. We therefore expect the latter to be local maximizers. This is confirmed by evaluating $m(c_1)$ and $m(c_2)$ as follows: for these values of c , we solve the system (5.33) with the last equation removed. This is a square system of equations⁷ that encodes Lagrange first order conditions for the Lie derivative, its solutions therefore include the global maximizer. The global maximizers obtained this way are shown in Table 5.2 and confirm that c_2 is indeed the unique positive root of $m(c)$.

Surprisingly, both $m(c)$ computed here and $\bar{m}(c)$ computed in the previous section have the same unique positive root (compare Table 5.1 and Table 5.2), while the upper bound (5.21) should entail some overestimation. This coincidence is left for future investigations.

5.2 Extended convergence condition for the super-twisting algorithm

In this section, we will consider a one-dimensional system

$$\dot{x}(t) = h(t) + g(t)u(t), \quad (5.34)$$

where $x(t)$ is typically the sliding variable of an affine controlled system, $g(t)$ and $h(t)$ being unknown with bounds discussed later. The sliding variable is to be driven to zero in finite time, and we consider the so-called super twisting algorithm (STA) defined by

$$u(t) = -k_1 \phi(x(t)) - k_2 \int_{\tau} \text{sign } x(\tau) d\tau, \quad (5.35)$$

7. With respect to the square system (5.33), the variable c is now fixed, and one equation has been removed, hence leading to a new square system.

where $\phi(x) = |x|^{\frac{1}{2}} \text{sign } x$. Introducing the new variables $x_1 = x$ and $x_2 = -k_2 \int_{\tau} \text{sign } x(\tau) d\tau$, the system (5.34) and its control (5.35) are rewritten as the two-dimensional right-hand side discontinuous ordinary differential equation

$$\dot{x}_1(t) = h(t) + g(t) \left(-k_1 \phi(x_1(t)) + x_2(t) \right) \quad (5.36a)$$

$$\dot{x}_2(t) = -k_2 \text{sign } x_1(t), \quad (5.36b)$$

whose solutions can be understood in the sense of Filippov [43]. The integral correction $x_2(t)$ does not converge to zero, therefore only the finite time convergence of $x_1(t)$ to zero is investigated.

As far as the authors know, the STA with such an uncertain input gain has only been investigated using majorant curves arguments, e.g., in [96]. Lyapunov functions represent another key approach for proving the finite time stability of such sliding mode controllers [76, 88, 93, 94], which offers several advantages. In addition to the simplicity of the argument once the correct Lyapunov function is discovered, they often offer some insights: e.g., the twisting controller can be assimilated to a nonsmooth harmonic oscillator with dissipation that enjoys an energy related Lyapunov function (see, e.g., Example 3 of [33, p.63]). The Lyapunov approach has been applied successfully to the STA by several authors [76, 88, 93] in the particular case where $g(t) = 1$.

Here, we build on the approach of Seeber and Horn [93], who have investigated the finite time stability of the STA subject to the following perturbation:

$$\dot{x}_1(t) = -k_1 \phi(x_1(t)) + x_2(t) \quad (5.37a)$$

$$\dot{x}_2(t) = -k_2 \text{sign } x_1(t) + \delta(t), \quad (5.37b)$$

where $\delta(t) \in [-L, L]$ is absolutely continuous. Finite-time stability of the STA (5.37) means converging to $x_1 = 0$ and $x_2 = 0$ in finite time. They provided the following sufficient conditions for finite time stability:

Theorem 5.3 ([93]). *The system (5.37) is finite time stable if $k_2 > L$ and $k_1^2 > k_2 + L$.*

Theorem 5.3 also applies to (5.36) with fixed input gain $g(t) = 1$ by noticing that the system (5.37) is equivalent to (5.34) and (5.35) with

$$h(t) = \delta_0 + \int_{\tau=0}^t \delta(\tau) d\tau. \quad (5.38)$$

This classical argument is to be used carefully in the context of Filippov solutions, the change of variable used in Section 5.2.2 formalizes it.

In this section, we derive from Theorem 5.3 new finite time stability sufficient conditions that applies to uncertain input gains in (5.36). We classically assume as in [96] that $g(t)$ and $h(t)$ satisfy $0 < \underline{g} \leq g(t) \leq \bar{g}$, $|h(t)| \leq \bar{h}$, and that they are continuously differentiable. As a consequence, $\gamma(t) = h(t)/g(t)$ is also continuously differentiable, and we furthermore assume $|\dot{\gamma}(t)| \leq M$. If $|\dot{g}(t)| \leq G$ and $|\dot{h}(t)| \leq H$, then $M = H/\underline{g} + G\bar{h}/\underline{g}^2$ is a valid bound for $|\dot{\gamma}(t)|$. Our main result is the following theorem, whose proof is given in Section 5.2.2:

Theorem 5.4. *The system (5.36) is finite time stable if $k_2 > M$ and $\underline{g}k_1^2 > k_2 + M$.*

The differentiability with bounded derivative is typically assumed for the investigation of the finite time convergence of the STA, e.g., in [88, 96]. In [76] Theorem 3 and in [93] Section 3, $g(t) = 1$ but $h(t)$ is not assumed to be bounded or differentiable, and finite time convergence is proven only under some technical additional assumption entailing that $h(t)$ vanishes at the origin⁸. Theorem 5.4 could be extended to handle the perturbation $\eta(t)\sqrt{|x_1|}$ used in [93] but details are not provided here for the sake of simplicity and clarity⁹. Dealing with unbounded uncertainties that are allowed to grow linearly with x_1 , typically satisfying $|h(t)| \leq \alpha_1 + \alpha_2|x_1|$, increases the practical scope of the investigation. Up to the knowledge of the authors, it requires modifications of the STA, e.g., generalizing the command law [15, 76]¹⁰. Other results on output feedback sliding mode controllers also achieve handling unbounded perturbations, e.g., [89]. Multivariable STA with uncertain input matrix is studied in [77], but it considers a generalized command law as well.

In Section 5.2.2 two consecutive changes of variables will transform (5.36) into (5.37) with a time-varying integral gain $k_2(t)$. The extension of stability properties to time variable gains is not trivial (see, e.g., Section 6 in [76]). The extension of Theorem 5.3 to time variable integral gain $k_2(t)$ is performed in Section 5.2.1. Finally, the conditions of Theorem 5.4 is compared to the classical conditions [88, 96] in Section 5.2.3.

8. The condition $|h(t)| \leq \eta(t)\sqrt{|x_1|}$ is used in [93] and the condition $|h(t)| \leq \eta(t)\sqrt{|x_1| + x_1^2}$ in [76]. Both allow unbounded $|h(t)|$ but entail $h(t)$ converges to zero when $x_1(t)$ does.

9. The constraint in Theorem 5.4 would then be $\underline{g}k_1^2 > k_2 + M + \eta_M$, where η_M would be a bound on $\eta(t)$.

10. The conditions in [15, 76] do not apply to the STA.

5.2.1 Extension to a time-varying integral gain

We consider here that the integral gain k_2 in system (5.37) is a time-varying gain $k_2(t)$, and we generalize Seeber and Horn’s sufficient conditions to accept a time-varying integral gain $k_2(t)$. Its proof is an extension of Seeber and Horn’s original proof: their argument is based on the (time-invariant) Lyapunov function,

$$v(\mathbf{x}) = \begin{cases} 2\sqrt{x_2^2 + 3\alpha^2 k_1^2 x_1} - x_2 & \text{if } \mathbf{x} \in \Sigma_1 \\ 2\sqrt{x_2^2 - 3\alpha^2 k_1^2 x_1} + x_2 & \text{if } \mathbf{x} \in \Sigma_2 \\ 3|x_2| & \text{otherwise,} \end{cases} \quad (5.39)$$

where $\Sigma_1 = x_1 \geq 0 \cap x_2 \leq \alpha k_1 \sqrt{x_1}$ and $\Sigma_2 = x_1 \leq 0 \cap x_2 \geq \alpha k_1 \sqrt{-x_1}$. The argument readily extends to a variable integral gain since k_2 does not influence the Lyapunov function, preserving the negativeness of the Lie derivative where it exists. However, their proof cannot be extended to time-varying gain $k_1(t)$ because (5.39) depends on k_1 .

We deal with the states where the locally Lipschitz Lyapunov function is not differentiable by using the nonsmooth version of Lyapunov’s stability theorem (see, e.g., [33]). Let $\tilde{\mathcal{L}}_f v(\mathbf{x})$ denote the set-valued Lie derivative of v w.r.t. f evaluated at \mathbf{x} , where f represents the system’s vector field. Intuitively, at the points where $\dot{v}(\mathbf{x})$ does not exist $\tilde{\mathcal{L}}_f v(\mathbf{x})$ is the convex hull of the limit points of neighbor Lie derivatives. We recall that, in analogy with Lyapunov’s stability theorem, one needs to investigate the negativeness of $\max \tilde{\mathcal{L}}_f v(\mathbf{x})$ to assess the system’s stability (see, e.g., [33] for details). Before stating Theorem 5.3, we introduce Lemma 5.1 which, when f is continuous, allows to assess the negativeness of $\tilde{\mathcal{L}}_f v(\mathbf{x})$ by considering only the values of $\dot{v}(\mathbf{x})$ where $v(\mathbf{x})$ is differentiable.

Lemma 5.1. *Consider an open subset \mathcal{N} of the state space. We assume that $f(t, \mathbf{x})$ is continuous w.r.t. \mathbf{x} for all $\mathbf{x} \in \mathcal{N}$ and all t , and that $v(\mathbf{x})$ be locally Lipschitz. Let $r(\mathbf{x})$ be a real valued function continuous on \mathcal{N} . If, for all t , $\dot{v}(t, \mathbf{x}) = \nabla v(\mathbf{x})^\top f(t, \mathbf{x}) \leq r(\mathbf{x})$ on every point where v is differentiable on \mathcal{N} , then, for all t , $\max \tilde{\mathcal{L}}_f v(t, \mathbf{x}) \leq r(\mathbf{x})$ on \mathcal{N} .*

Proof. Let t be an arbitrary time instant and let Ω_v be the set of non-differentiability points of v inside \mathcal{N} . On $\mathcal{N} \setminus \Omega_v$, $\tilde{\mathcal{L}}_f v(t, \mathbf{x}) = \{\dot{v}(t, \mathbf{x})\}$, so the lemma is trivially true. For $\mathbf{x} \in \Omega_v$, by the continuity of f we have that $\tilde{\mathcal{L}}_f v(t, \mathbf{x})$ consists of the convex combinations of limit points of $\dot{v}(t, \mathbf{x}_i)$ at neighboring points \mathbf{x}_i where v is differentiable (this is a direct consequence of the definition of $\tilde{\mathcal{L}}_f v$ and a few of its basic properties which can be found, e.g., in [33]). Since $\tilde{\mathcal{L}}_f v(t, \mathbf{x})$ is a compact interval of real numbers, its max is a vertex of

the one dimensional convex hull, meaning that there exists a converging sequence $\mathbf{x}_i \rightarrow \mathbf{x}$ such that $\dot{v}(t, \mathbf{x}_i) \rightarrow \max \tilde{\mathcal{L}}_f v(t, \mathbf{x})$. But $\dot{v}(t, \mathbf{x}_i) \leq r(\mathbf{x}_i)$ for every point in the sequence, meaning that, by continuity of r , we have $\max \tilde{\mathcal{L}}_f v(t, \mathbf{x}) \leq r(\mathbf{x})$, which holds for all times since t is arbitrary. \square

Remark. Lemma 5.1 allows for any function $r(\mathbf{x})$ to be used, in practice $r(\mathbf{x}) = -\epsilon$ or $r(\mathbf{x}) = -v(\mathbf{x})^\alpha$, $0 < \alpha < 1$, can be used to assess finite time stability.

The following theorem extends Seeber and Horn's theorem to a time-varying integral gain. The assumption that $k_2(t)$ is strictly greater than $|\delta(t)|$ can be stated in two natural ways: first by using a lower bound $\underline{k}_2 \leq k_2(t)$ and enforcing $\underline{k}_2 \geq L$ as in Theorem 5.3. Second, by enforcing the inequality $k_2(t) \geq |\delta(t)| + \epsilon$ at each time, with an arbitrary $\epsilon > 0$. The first condition is sensibly stronger than the second, and leads to a more restrictive finite time stability condition than Theorem 5.4 (with an additional factor \bar{g}/\underline{g}). This is why the second statement is used in the following time-varying integral gain version of Theorem 5.3:

Theorem 5.5. *The system (5.37) with time-varying integral gain $k_2(t)$ is finite time stable if there exists $\epsilon > 0$ such that $k_2(t) \geq |\delta(t)| + \epsilon$ and $k_1^2 \geq k_2(t) + |\delta(t)| + \epsilon$ holds for all times.*

Proof. The same Lyapunov function (5.39) is used here, except for the choice of the parameter $\alpha < 1$ that needs to account for ϵ in the statement: here one can choose α such that

$$(\alpha k_1)^2 \geq k_2(t) + |\delta(t)| + \frac{1}{2}\epsilon \quad (5.40)$$

Wherever the Lyapunov function is differentiable, the upper bound on the Lie derivative computed in (14) in [93] still holds when k_2 varies in time, meaning that

$$\dot{v} \leq 3 \max \left\{ k_2(t) - \delta(t) - (\alpha k_1)^2, \frac{1}{2}\alpha(\alpha - 1)k_1^2 \right\}. \quad (5.41)$$

As in the original proof, (5.40) entails $k_2(t) - \delta(t) - (\alpha k_1)^2 \leq -\frac{1}{2}\epsilon$ and $0 < \alpha < 1$ entails $\alpha(\alpha - 1) < 0$, showing that the maximum of them is strictly negative.

Now, we need to investigate the states where v is not differentiable. It is well known that no sliding motion can occur along the set $S_1 = \{\mathbf{x} | x_1 = 0, x_2 \neq 0\}$, whether k_2 is fixed or time-varying. The other set where v is not differentiable is $S_2 = \{\mathbf{x} | x_2 = \alpha k_1 \sqrt{|x_1|}, x_1 \neq 0\}$. Since system (5.37) is continuous on $\mathcal{N} := \mathbb{R}^2 \setminus S_1$, we can apply

Lemma 5.1 with $r(\mathbf{x}) = 3 \max \left\{ -\frac{1}{2}\epsilon, \frac{1}{2}\alpha(\alpha - 1)k_1^2 \right\}$ to show that the same strictly negative upper bound that holds when v is differentiable also holds on the whole set \mathcal{N} , which includes the set S_2 . \square

5.2.2 Extension to uncertain input gain

The STA (5.36) aims $x_1(t) = 0$ and $\dot{x}_1(t) = 0$, the integral correction $x_2(t)$ being foreseen not to converge. Indeed, the value of $x_2(t)$ enforced at steady state can be computed with respect to uncertainties: if $x_1(t) = 0$ and $\dot{x}_1(t) = 0$ then $x_2(t) = -h(t)/g(t)$. Classically, by defining the system $y_1(t) = x_1(t)$ and $y_2(t) = x_2(t) + h(t)/g(t)$, we obtain a system that is expected to converge to zero. It satisfies the following ODE:

$$\dot{y}_1(t) = g(t) \left(-k_1 \phi(y_1(t)) + y_2(t) \right) \quad (5.42a)$$

$$\dot{y}_2(t) = -k_2 \operatorname{sign} y_1(t) + \dot{\gamma}(t), \quad (5.42b)$$

with

$$\dot{\gamma}(t) = \frac{d}{dt} \frac{h(t)}{g(t)} = \frac{\dot{h}(t)}{g(t)} - \frac{\dot{g}(t) h(t)}{g(t)^2}. \quad (5.43)$$

The change of variable $\mathbf{x} \mapsto \mathbf{y}$ is a diffeomorphism for all times, therefore Theorem 2 of [43, p.99] proves that Filippov solutions $\mathbf{x}(t)$ map to Filippov solutions $\mathbf{y}(t)$ and vice-versa. Since $y_1(t) = x_1(t)$ the finite time stability of the (5.42) is equivalent to the finite time stability of (5.36).

We introduce a rescaling of time $\mathbf{z}(t) = \mathbf{y}(\tau(t))$, that is going to move the impact of the perturbation $g(t)$ from (5.42a) to (5.42b), so as to match the perturbation structure of (5.37). We define $\tau(t)$ by the following ODE:

$$\dot{\tau}(t) = \frac{1}{g(\tau(t))}, \quad (5.44)$$

with $\tau(0) = 0$. Since $g(t) \in [\underline{g}, \bar{g}]$ is bounded away from zero by $\underline{g} > 0$ and is differentiable with bounded derivative $\dot{g}(t) \in [-G, G]$, the function $1/g(\tau)$ is continuously differentiable with bounded derivative. As a consequence the ODE (5.44) has a unique solution, which is globally defined in \mathbb{R} . Furthermore, from (5.44) and the bounds on $g(t)$ we obtain $1/\bar{g} \leq \dot{\tau}(t) \leq 1/\underline{g}$, which proves that $\tau(t)$ is strictly increasing and satisfies,

$$\frac{t}{\bar{g}} \leq \tau(t) \leq \frac{t}{\underline{g}}. \quad (5.45)$$

This makes $\tau(t)$ a valid rescale of time. The time-rescaled system $\mathbf{z}(t) = \mathbf{y}(\tau(t))$ satisfies the following relation:

$$\dot{\mathbf{z}}(t) = \dot{\tau}(t) \dot{\mathbf{y}}(\tau(t)) = \frac{1}{g(\tau(t))} \dot{\mathbf{y}}(\tau(t)), \quad (5.46)$$

obtained using the chain rule and (5.44). Finally using (5.42) we obtain the ODE governing $\mathbf{z}(t)$:

$$\dot{z}_1(t) = -k_1 \phi(z_1(t)) + z_2(t) \quad (5.47a)$$

$$\dot{z}_2(t) = -\frac{k_2}{\tilde{g}(t)} \text{sign } z_1(t) + \frac{\tilde{\gamma}(t)}{\tilde{g}(t)}, \quad (5.47b)$$

where $\tilde{g}(t) = g(\tau(t))$ and $\tilde{\gamma}(t) = \dot{\gamma}(\tau(t))$. The system (5.47) is in the scope of Theorem 5.5 with

$$k_2(t) = \frac{k_2}{\tilde{g}(t)} \quad \text{and} \quad \delta(t) = \frac{\tilde{\gamma}(t)}{\tilde{g}(t)}. \quad (5.48)$$

Its the finite time stability is equivalent to the finite time stability of (5.42) because of (5.45). The bounds $g(t) \in [\underline{g}, \bar{g}]$ and $\dot{\gamma}(t) \in [-M, M]$ hold for their time-shifted counterparts $\tilde{g}(t)$ and $\tilde{\gamma}(t)$.

We are finally ready to state the proof of Theorem 5.4.

Proof of Theorem 5.4. Suppose the conditions $k_2 > M$ and $\underline{g} k_1^2 > k_2 + M$ of Theorem 5.4 hold. Since $\tilde{\gamma}(t) \in [-M, M]$ and $\tilde{g}(t) > \underline{g}$, there exists $\epsilon > 0$ such that both $k_2 \geq |\tilde{\gamma}(t)| + \epsilon$ and $\tilde{g}(t) k_1^2 \geq k_2 + |\tilde{\gamma}(t)| + \epsilon$ hold for all times. Dividing both inequalities by $\tilde{g}(t) < \bar{g}$, we prove that

$$k_2(t) \geq |\delta(t)| + \frac{\epsilon}{\bar{g}} \quad \text{and} \quad k_1^2 \geq k_2(t) + |\delta(t)| + \frac{\epsilon}{\bar{g}}, \quad (5.49)$$

where $k_2(t)$ and $\delta(t)$ are defined in (5.48). This allows applying Theorem 5.5 to the STA (5.47) with (5.48), hence proving its finite time stability. \square

5.2.3 Comparison to the state-of-the-art conditions

In order to handle the same differentiable perturbations $g(t)$ and $h(t)$ as in (5.36), Shtessel, Edwards, Fridman and Levant [96] investigate a modified STA, where a switching law is added to keep u bounded in $[-U_m, U_m]$ and to prove the finite time convergence. They proved some sufficient conditions on k_1 , k_2 and U_m for the finite time stability,

see [96, p.156]. After some algebraic manipulations¹¹, one can prove that these conditions imply

$$k_2 > M \quad \text{and} \quad \underline{g} k_1^2 > 2 \left(\frac{\bar{g}}{\underline{g}} \right)^2 (k_2 + M), \quad (5.50)$$

where $M = H/\underline{g} + G \hbar/\underline{g}^2$. The first bound on k_2 is the same as in Theorem 5.4. However, the second bound on k_1^2 is seen to be similar but more restrictive than the one in Theorem 5.4 because of the factor $2(\bar{g}/\underline{g})^2$. In particular, if the interval $[\underline{g}, \bar{g}]$ is large and/or \underline{g} is close to zero then $(\bar{g}/\underline{g})^2$ is sensibly greater than 1 and the conditions of [96] become sensibly more restrictive than Theorem 5.1.

11. Among other basic manipulations, setting q and U_m to the best cases $q \in \{0, 1\}$ and $U_m = C$ in order to weaken the conditions, and using the fact that $0 < \underline{g}k_2 - C \leq \underline{g}k_2 + C$ with $C = H + G \hbar/\underline{g} = \underline{g} M$ being the smallest upper bound to $|\dot{h}(t) + \dot{g}(t) u(t)|$ subject to $|u(t)| \leq U_m$ that can be obtained using bounds on the perturbations provided above.

CONCLUSION AND PERSPECTIVES

In this thesis, we deeply explored the stability properties of image-based visual servoing systems, both by developing new theoretical tools that allow us to study the stability of *energy-decreasing* systems and by creating a new, *ad hoc* strategy that makes the computation of the full set of equilibria tractable in the context of visual servoing. Thanks to our new methodology, which is inspired by Lyapunov and Morse theories, we were able to rigorously characterize the interconnection between the equilibria of these systems and to propose the *controller-independent* regions of attraction of the stable ones. Moreover, while the main objective was the analysis of image-based visual servoing systems, the tools that we developed allowed us to synthesize a new control paradigm, that we call *meta-controller*, which provides a substantial improvement in size for the desired pose's region of attraction, showcasing the importance of understanding the stability of these systems and the importance of saddles in their analysis.

Overall, our strategy grants us the ability to analyze a large class of classical situations in a more profound way, shining a new light on the behavior of these systems. In the authors knowledge, this is the first time that all stable and unstable equilibria of these systems have been identified, which allows us to provide a deeper qualitative understanding on the system's dynamics. Moreover, it opens the door to a rich set of previously unknown behaviors, as well as providing a new way to study known phenomena in the field. Hopefully, this might be just the beginning of a line of research that aims to build up a global understanding of the dynamics of these systems.

From a computational point of view, we believe that, as the solving algorithms and hardware evolve, attacking examples with configurations of 5 points will be within reach in the near future. Other strategies could significantly improve the performances that we obtained with `msolve`: for instance, it is possible to show that the equilibrium condition that we derived back in Chapter 2 is equivalent to a non-linear optimization problem with equality constraints, which means that we can reformulate the equilibrium condition using the Lagrange multipliers formalism. Other non-trivial changes of variables, on the same spirit of the one that we proposed in Section 2.3.2, could also lower the computation time.

A different approach would be to change the class of solvers that we use. In the past, we briefly explored the possibility of using numerical solvers based on homotopy continuation, which, unfortunately, are not guaranteed to find all solutions in a rigorous way (see [45] for a comparison between `msolve` and these methods); however, developing our own solver that rigorously perform the continuation is possible, in particular by employing interval analysis-based methods.

Solving the equilibrium condition directly with interval analysis based solvers, like IBEX, can also become a viable strategy if we can find a reasonable bounded initial domain that is guaranteed to contain all solutions. In this regard, we had some promising preliminary results by *homogenizing* the equilibrium condition’s system of equation, which is a classical technique that allows to project the set of solutions of an n -dimensional system of polynomials onto a sphere in \mathbb{R}^{n+1} , effectively *compactifying* the system’s domain. These results, among others, were studied by Mattia Piras in his Master’s thesis [87], carried out under our supervision.

Since, as we hinted in Section 1.1, the characterization of image-based visual servoing controllers does not depend on the fact that we track the points’ Cartesian coordinates, another interesting generalization would be a different choice of visual features, like the polar coordinates of the points in the image, the lines or even the image moments. Each of these choices leads to wildly different interaction matrices and, thus, equilibrium conditions; however, we believe that the methodology that we propose of switching over to the (extended) features space, together with the addition of new constraints to enforce the feasibility of solutions, could represent a potential game changer also in these cases.

A last interesting line of research, for what concerns the more practical aspects of our contributions, would be including the overall robotic system (i.e., the camera together with the robotic arm or other “supports”) in the investigation of the stability. Not only this would bring our results closer to real-world applications, but it would also allow us to test our findings on real robotics systems. This is particularly interesting for validating whether the more theoretical among our results do hold in practice, like, for instance, if we can verify the presence of saddle points or the region of attraction’s enlargement for the meta-controller.

The theoretical contributions of our thesis also open up several interesting questions, some more technical than the others. As it is often the case when we study the properties of a class of systems, we immediately feel the need to try and admit more and more systems within the class, and there is a number of generalizations that feel natural in our case. The

first restriction that we encountered in Chapter 1 was leaving out of the discussion stable space of non-differentiable systems. However, we did present some examples where the stable spaces seemed to make sense, simply losing its smoothness. Thus, it seems possible to identify some additional hypotheses, so that we can discuss the stable and unstable manifolds of a specific class of non-differentiable systems.

Another possible generalization could be relaxing the definition of energy-decreasing systems, allowing the Lie derivative to be null outside of the energy function's critical points, as long as the set of states for which this happens is not invariant. With this relaxation, we could consider extend our analysis, for instance, to dissipative physical systems whose energy does not satisfy the Lyapunov condition globally, a classical example being the pendulum with friction (whose Lie derivative is null at every peak of the swing).

At the current state, the construction of the connectivity tree only makes sense as long as the sublevel sets are compact. However, we believe that, if we have unbounded level sets in \mathbb{R}^n , it is reasonable to model the sublevel set becoming unbounded as it “reaching” a critical point which is arbitrarily far. Consider, for instance, the classical example we presented in Figure 1.13 (Section 1.4): if we restrict on the x_1 -axis, as we let x_1 grow indefinitely, the function will increase up to a maximal, bounded value. In turn, the same value is a local minimum along the x_2 direction, so it seems intuitive to say that, as $x_1 \rightarrow \infty$ with $x_2 = 0$, we are converging towards a *saddle point*. We believe that these *critical points at infinity* could be identified and studied by using the techniques that we discussed above to compactify the space, and their characterization could allow to push the connectivity tree further up in these cases.

A final topic which has not been completely fleshed out is the separation of the sublevel sets' components. Back in Section 1.3.5, in order to make sense of the separating hyperplanes we had to assume that our state-space was Euclidean. While we have little hope to be able to find a generalization for the separating function that works on any arbitrary manifold, it would be interesting to find such a generalization for some manifolds of interest in robotics, like SE(3) or SE(2).

PROOF OF THEOREM 1.1

Before diving into the proof, we recall two useful matrix properties and we state a lemma that we will need for the theorem.

Remark. Consider $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{n \times n}$. We have that:

- If they are similar (i.e., there exists an invertible matrix \mathbf{P} such that $\mathbf{A} = \mathbf{P}^{-1}\mathbf{B}\mathbf{P}$), then they have the same eigenvalues.
- If they are symmetric and congruent (i.e., there exists an invertible matrix \mathbf{P} such that $\mathbf{A} = \mathbf{P}^\top \mathbf{B} \mathbf{P}$), then, by Sylvester's law of inertia, they have the same signature.

Lemma A.1. Let $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{n \times n}$ be symmetric matrices and let \mathbf{A} be positive definite as well. Then, $\mathbf{A}\mathbf{B}$ has real eigenvalues and the same signature as \mathbf{B} .

Proof of Lemma A.1. Being \mathbf{A} symmetric positive definite (SPD), there exists an SPD matrix $\mathbf{A}^{\frac{1}{2}}$ such that $\mathbf{A}^{\frac{1}{2}}\mathbf{A}^{\frac{1}{2}} = \mathbf{A}$. By defining $\mathbf{C} = \mathbf{A}^{\frac{1}{2}}\mathbf{B}\mathbf{A}^{\frac{1}{2}}$, we have that:

$$\mathbf{A}^{\frac{1}{2}}\mathbf{C}\mathbf{A}^{-\frac{1}{2}} = \mathbf{A}^{\frac{1}{2}}\mathbf{A}^{\frac{1}{2}}\mathbf{B}\mathbf{A}^{\frac{1}{2}}\mathbf{A}^{-\frac{1}{2}} = \mathbf{A}\mathbf{B}, \quad (\text{A.1})$$

meaning that \mathbf{C} and $\mathbf{A}\mathbf{B}$ are similar. But since $\mathbf{A}^{\frac{1}{2}}$ is symmetric, \mathbf{C} is symmetric and it can be written as $\mathbf{C} = (\mathbf{A}^{\frac{1}{2}})^\top \mathbf{B} \mathbf{A}^{\frac{1}{2}}$, meaning that \mathbf{C} and \mathbf{B} are congruent. Having $\mathbf{A}\mathbf{B}$ the same eigenvalues of \mathbf{C} , this means that they are real and that $\mathbf{A}\mathbf{B}$ and \mathbf{B} have the same signature. \square

We are ready to prove the theorem.

Proof of Theorem 1.1. We argue that all gradient systems on the same energy function have the same index as the system defined using the standard gradient, for any given equilibrium. Let $v : \mathcal{M} \rightarrow \mathbb{R}$ be the energy function and $\mathbf{x} \in \mathcal{M}$ be an equilibrium. The choice of metric on the tangent space $T_{\mathbf{x}}\mathcal{M}$ induces a map $\Phi : T_{\mathbf{x}}\mathcal{M} \rightarrow T_{\mathbf{x}}^*\mathcal{M}$ from the tangent space itself to its dual by:

$$\Phi(\mathbf{u})(\mathbf{v}) = \langle \mathbf{u}, \mathbf{v} \rangle_{\mathbf{x}}, \quad (\text{A.2})$$

where $\mathbf{u}, \mathbf{v} \in T_{\mathbf{x}}\mathcal{M}$. This map is an isomorphism, and as such it is invertible. Endowing the tangent space and its dual with their canonical bases and identifying the inner product with an SPD matrix \mathbf{P} as before, we can find explicit expression for Φ and its inverse as:

$$\begin{cases} \Phi(\mathbf{u}) = \mathbf{u}^\top \mathbf{P} \\ \Phi^{-1}(\mathbf{u}^\top) = \mathbf{P}^{-1} \mathbf{u} \end{cases}, \quad (\text{A.3})$$

where we denote elements of $T_{\mathbf{x}}\mathcal{M}$ as column vectors and elements of its dual as row vectors to ease the notation.

We need now to be careful with how to discuss the Hessian of v , because, among the several equivalent ways found in the literature, we want to use a formalism that highlights the role of the metric. Let $D_{\mathbf{x}}^2 v : T_{\mathbf{x}}\mathcal{M} \times T_{\mathbf{x}}\mathcal{M} \rightarrow \mathbb{R}$ be the usual second derivative of v computed at \mathbf{x} , which is a symmetric bilinear map independent from the metric. Based on it, we can define two flavors of the Hessian. We call the Hessian *map* of v at \mathbf{x} the linear map $\mathcal{H}_{\mathbf{x}} : T_{\mathbf{x}}\mathcal{M} \rightarrow T_{\mathbf{x}}^*\mathcal{M}$ defined by¹:

$$\mathcal{H}_{\mathbf{x}}(\mathbf{u})(\mathbf{v}) = D_{\mathbf{x}}^2 v(\mathbf{u}, \mathbf{v}). \quad (\text{A.4})$$

The Hessian map depends on the second derivative alone, and thus it is not influenced by the metric. Building upon it, we define the Hessian *matrix* of v at \mathbf{x} the *endomorphism*² $\mathbf{H}_{\mathbf{x}} = \Phi^{-1} \circ \mathcal{H}_{\mathbf{x}} : T_{\mathbf{x}}\mathcal{M} \rightarrow T_{\mathbf{x}}\mathcal{M}$, which now clearly depends on the metric through Φ . If we denote with $\Phi_{\mathbf{I}}$ the mapping induced by the standard inner product, we get that $\mathbf{H}_{\mathbf{x}, \mathbf{I}} = \Phi_{\mathbf{I}}^{-1} \circ \mathcal{H}_{\mathbf{x}}$ corresponds to the standard Hessian matrix, whose matrix representation is then symmetric.

By choosing bases for the tangent space and its dual, and by committing the minor sin of identifying endomorphisms with their matrix representations, we obtain the following relation:

$$\mathbf{H}_{\mathbf{x}} = \Phi^{-1} \circ \mathcal{H}_{\mathbf{x}} = \Phi^{-1} \circ \Phi_{\mathbf{I}} \circ \Phi_{\mathbf{I}}^{-1} \circ \mathcal{H}_{\mathbf{x}} = (\Phi^{-1} \circ \Phi_{\mathbf{I}}) \mathbf{H}_{\mathbf{x}, \mathbf{I}}. \quad (\text{A.5})$$

We can derive the matrix representation of the endomorphism $\Phi^{-1} \circ \Phi_{\mathbf{I}}$ by using (A.3).

1. The identification of bilinear maps from a vector space to the reals with linear maps from a vector space to its dual is actually a simple application of a much more general property: if $L(E_1, \dots, E_n; F)$ denotes the set of n -linear maps from $E_1 \times \dots \times E_n$ to F , then there is a natural isomorphism between $L(E_1, \dots, E_{n+k}; F)$ and $L(E_1, \dots, E_n; L(E_{n+1}, \dots, E_{n+k}; F))$, even when $E_1 \times \dots \times E_n$ and F are infinite-dimensional (see, e.g., [62, Proposition 5.2]).

2. I.e., a linear map from a vector space to itself.

Letting \mathbf{P} denote the SPD matrix representing the inner product related to Φ , we have:

$$(\Phi^{-1} \circ \Phi_{\mathbf{I}})(\mathbf{u}) = \Phi^{-1}(\mathbf{u}^{\top}) = \mathbf{P}^{-1}\mathbf{u}, \quad (\text{A.6})$$

which leads to:

$$\mathbf{H}_{\mathbf{x}} = \mathbf{P}^{-1}\mathbf{H}_{\mathbf{x},\mathbf{I}}. \quad (\text{A.7})$$

We can finally apply Lemma A.1, proving that $\mathbf{H}_{\mathbf{x}}$ and $\mathbf{H}_{\mathbf{x},\mathbf{I}}$ have the same signature and, thus, that the index of \mathbf{x} is the same for all gradient systems. \square

Remark. *By applying Lemma A.1, we actually prove something stronger: even though $\mathbf{H}_{\mathbf{x}}$ might not be symmetric for a given metric, it must always have real eigenvalues, which justifies Corollary 1.1.*

BIBLIOGRAPHY

Author's publications

- [29] A. Colotti and A. Goldsztejn, « Practical stability and attractors of systems with bounded perturbations », *in: 2022 IEEE 61st Conference on Decision and Control (CDC)*, IEEE, 2022, pp. 5129–5134.
- [32] A. Colotti, D. Monnet, A. Goldsztejn, and F. Plestan, « New convergence conditions for the super twisting algorithm with uncertain input gain », *in: Automatica* 143 (2022), p. 110423.
- [45] J. Garcia Fontan, A. Colotti, S. Briot, A. Goldsztejn, and M. Safey El Din, « Computer Algebra Methods for Polynomial System Solving at the Service of Image-Based Visual Servoing », *in: ACM Commun. Comput. Algebra* 56.2 (Nov. 2022), pp. 36–40.

In review

- [30] A. Colotti, C. Lu, A. Goldsztejn, O. Kermorgant, and Y. Li, « An energy-decreasing controller for sensor-based control with enhanced region of attraction », *in: Submitted to Robotics and Automation Letters* (2023).
- [31] A. Colotti, J. Garcia Fontan, A. Goldsztejn, S. Briot, F. Chaumette, O. Kermorgant, and M. Safey El Din, « Determination of All Stable and Unstable Equilibria for Image Points-Based Visual Servoing », *in: Submitted to IEEE Trans. on Robotics* (2023).

Bibliography

- [1] P.-A. Absil, R. Mahony, and B. Andrews, « Convergence of iterates of descent methods for analytic cost functions », *in: SIAM Journal on Optimization* (2005).
- [2] J. Alexander and J. A. Yorke, « The homotopy continuation method: numerically implementable topological procedures », *in: Trans. of the American Mathematical Society* 242 (1978), pp. 271–284.
- [3] C. Aliprantis, *Infinite dimensional analysis : a hitchhiker's guide*, Berlin New York: Springer, 2006, ISBN: 9783540295860.
- [4] T. M. Apostol, *Calculus, Volume 2*, John Wiley & Sons, 1991.
- [5] I. Araya et al., « Upper Bounding in Inner Regions for Global Optimization Under Inequality Constraints », *in: J. Global Optimization* 60.2 (2014), pp. 145–164.
- [6] J.-P. Aubin and H. Frankowska, *Set-Valued Analysis*, Birkhäuser Boston, 2009, DOI: 10.1007/978-0-8176-4848-0.
- [7] B. Bank et al., *Non-Linear Parametric Optimization*, Birkhäuser Basel, 1982, DOI: 10.1007/978-3-0348-6328-5.
- [8] A. Banyaga and D. Hurtubise, *Lectures on Morse Homology*, Springer Netherlands, 2004, DOI: 10.1007/978-1-4020-2696-6.
- [9] J. Berthomieu, C. Eder, and M. Safey El Din, « msolve: A Library for Solving Polynomial Systems », *in: ACM Int. Symp. on Symbolic and Algebraic Computation (ISSAC'21)*, Saint Petersburg, Russia, July 2021, pp. 51–58, DOI: 10.1145/3452143.3465545.
- [10] F. Blanchini, « Set invariance in control », *in: Automatica* 35.11 (Nov. 1999), pp. 1747–1767, DOI: 10.1016/s0005-1098(99)00113-2.
- [11] K. C. Border, « Introduction to Correspondences », Lecture notes, from: healy.econ.ohio-state.edu/kcb/Notes/Correspondences.pdf. California Institute of Technology, 2016.
- [12] B. Buchberger and F. Winkler, *Gröbner bases and applications*, vol. 17, Cambridge University Press Cambridge, 1998.
- [13] S. L. Campbell and C. D. Meyer, *Generalized Inverses of Linear Transformations*, Cambridge, Mar. 2009, 184 pp., ISBN: 0898716713.

-
- [14] J. F. Canny, E. Kaltofen, and L. Yagati, « Solving systems of nonlinear polynomial equations faster », *in: ACM-SIGSAM Int. Symp. on Symbolic and Algebraic Computation (ISSAC'89)*, 1989, pp. 121–128.
- [15] I. Castillo, L. Fridman, and J. A. Moreno, « Super-Twisting Algorithm in presence of time and state dependent perturbations », *in: International Journal of Control* 91.11 (2018), pp. 2535–2548, DOI: 10.1080/00207179.2016.1269952.
- [16] E. Cervera et al., « Improving image-based visual servoing with three-dimensional features », *in: The Int. Journal of Robotics Research* 22.10-11 (2003), pp. 821–839.
- [17] G. Chabert and L. Jaulin, « Contractor Programming », *in: Artif. Intell.* 173.11 (2009), pp. 1079–1100.
- [18] G. Chabert, *IBEX: a C++ library for constraint processing over real numbers*, <https://github.com/ibex-team/ibex-lib>, 2011.
- [19] F. Chaumette, « Image Moments: A General and Useful Set of Features for Visual Servoing », *in: IEEE Transactions on Robotics* 20.4 (Aug. 2004), pp. 713–723, DOI: 10.1109/tro.2004.829463.
- [20] F. Chaumette, « Potential problems of stability and convergence in image-based and position-based visual servoing », en, *in: The confluence of vision and control*, ed. by D. J. Kriegman, G. D. Hager, and A. S. Morse, Lecture Notes in Control and Information Sciences, London: Springer, 1998, pp. 66–78, ISBN: 9781846285288, DOI: 10.1007/BFb0109663.
- [21] F. Chaumette and E. Malis, « 2 1/2 D visual servoing: a possible solution to improve image-based and position-based visual servoings », *in: IEEE Int. Conf. on Robotics and Automation, ICRA'00*, San Francisco, CA, Apr. 2000, pp. 630–635.
- [22] F. Chaumette, « Potential problems of stability and convergence in image-based and position-based visual servoing », *in: The confluence of vision and control*, Springer, 2007, pp. 66–78.
- [23] F. Chaumette and S. Hutchinson, « Visual Servo Control, Part I: Basic Approaches », *in: IEEE Robotics & Automation Magazine* 13.4 (Dec. 2006), pp. 82–90, DOI: 10.1109/mra.2006.250573.

-
- [24] C. C. Cheah, S. Kawamura, and S. Arimoto, « Feedback control for robotic manipulator with uncertain kinematics and dynamic », *in: IEEE International Conference on Robotics and Automation*, 1998.
- [25] G. Chesi, *Domain of Attraction: Analysis and Control Via SOS Programming*, Springer-Verlag GmbH, Oct. 2011, ISBN: 0857299581.
- [26] G. Chesi et al., « Solving quadratic distance problems: an LMI-based approach », *in: IEEE Transactions on Automatic Control* 48.2 (Feb. 2003), pp. 200–212, DOI: 10.1109/tac.2002.808465.
- [27] H.-D. Chiang, M. Hirsch, and F. Wu, « Stability regions of nonlinear autonomous dynamical systems », *in: IEEE Trans. on Automatic Control* 33.1 (1988), pp. 16–27, DOI: 10.1109/9.357.
- [28] C. Collewet, E. Marchand, and F. Chaumette, « Visual servoing set free from image processing », *in: 2008 IEEE International Conference on Robotics and Automation*, IEEE, May 2008, DOI: 10.1109/robot.2008.4543190.
- [33] J. Cortes, « Discontinuous Dynamical Systems: A tutorial on solutions, nonsmooth analysis, and stability », *in: IEEE Control Systems Magazine* 28 (3) (2008), 36–73 (Jan. 2009), DOI: 10.1109/MCS.2008.919306, arXiv: 0901.3583.
- [34] M. Costanzo et al., « Stability and convergence analysis of 3D feature-based visual servoing », *in: IEEE Robotics and Automation Letters* 7.4 (Oct. 2022), pp. 12022–12029.
- [35] D. Cox, J. Little, and D. O’Shea, *Ideals, varieties, and algorithms: an introduction to computational algebraic geometry and commutative algebra*, Springer Science & Business Media, 2013.
- [36] H. B. Curry, « The method of steepest descent for non-linear minimization problems », *in: Quarterly of Applied Mathematics* (1944).
- [37] L. Deng, F. Janabi-Sharifi, and W. J. Wilson, « Stability and robustness of visual servoing methods », *in: IEEE International Conference on Robotics and Automation*, 2002.
- [38] B. Espiau, « Effect of camera calibration errors on visual servoing robotics », *in: 3rd International Symposium on Experimental Robotics*, Oct. 1993.
- [39] J.-C. Faugère, « A new efficient algorithm for computing Gröbner bases (F4) », *in: Journal of Pure and Applied Algebra* 139.1-3 (1999), pp. 61–88.

-
- [40] J.-C. Faugère, « A new efficient algorithm for computing Gröbner bases without reduction to zero (F5) », *in: ACM Int. Symposium on Symbolic and Algebraic Computation (ISSAC'02)*, 2002, pp. 75–83.
- [41] J.-C. Faugère, « FGb: a library for computing Gröbner bases », *in: Int. Congress on Mathematical Software*, Springer, 2010, pp. 84–87.
- [42] J.-C. Faugère et al., « Efficient computation of zero-dimensional Gröbner bases by change of ordering », *in: Journal of Symbolic Computation* 16.4 (1993), pp. 329–344.
- [43] A. F. Filippov, *Differential Equations with Discontinuous Righthand Sides*, Springer Netherlands, Sept. 1988, 316 pp., ISBN: 90-277-2699-X.
- [44] M. Fischler and R. Bolles, « Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography », *in: Communications of the ACM* 24.6 (June 1981), pp. 381–395.
- [46] P. Glendinning, *Stability, Instability and Chaos: an introduction to the theory of nonlinear differential equations*, Cambridge University Press, 1994.
- [47] A. Goldsztejn and L. Granvilliers, « A New Framework for Sharp and Efficient Resolution of NCSP with Manifolds of Solutions », *in: Constraints* 15.2 (2010), ed. by Springer, pp. 190–212, ISSN: 1383-7133.
- [48] A. Goldsztejn and G. Chabert, « Estimating the robust domain of attraction for non-smooth systems using an interval Lyapunov equation », *in: Automatica* 100 (Feb. 2019), pp. 371–377, DOI: 10.1016/j.automatica.2018.03.036.
- [49] G. Golub and C. Van Loan, *Matrix Computations*, J. Hopkins Uni. Press, Jan. 2013, ISBN: 1421407949.
- [50] J. C. Gower and G. B. Dijkstra, *Procrustes Problems*, Oxford University Press, 2004.
- [51] O. Hachicho and B. Tibken, « Estimating domains of attraction of a class of nonlinear dynamical systems with LMI methods based on the theory of moments », *in: Proceedings of the 41st IEEE Conference on Decision and Control, 2002*. IEEE, 2002, DOI: 10.1109/cdc.2002.1184354.
- [52] S. Hafstein and P. Giesl, « Review on computational methods for Lyapunov functions », *in: Discrete and Continuous Dynamical Systems - Series B* 20.8 (Aug. 2015), pp. 2291–2331, DOI: 10.3934/dcdsb.2015.20.2291.

-
- [53] M. W. Hirsch, R. L. Devaney, and S. Smale, *Differential Equations, Dynamical Systems, and Linear Algebra*, Elsevier Science & Techn., June 1974, 358 pp., ISBN: 9780080873763.
- [54] S. Hutchinson, G. D. Hager, and P. I. Corke, « A tutorial on visual servo control », *in: IEEE Transactions on Robotics and Automation* 12.5 (1996), pp. 651–670.
- [55] A. Iannelli, A. Marcos, and M. Lowenberg, « Robust estimations of the Region of Attraction using invariant sets », *in: Journal of the Franklin Institute* 356.8 (May 2019), pp. 4622–4647, DOI: 10.1016/j.jfranklin.2019.02.013.
- [56] W. Kabsch, « A discussion of the solution for the best rotation to relate two sets of vectors », *in: Acta Crystallographica Section A* 34.5 (Sept. 1978), pp. 827–828, DOI: 10.1107/s0567739478001680.
- [57] W. Kabsch, « A solution for the best rotation to relate two sets of vectors », *in: Acta Crystallographica Section A* 32.5 (Sept. 1976), pp. 922–923, DOI: 10.1107/s0567739476001873.
- [58] O. Kermorgant, *log2plot*, version 3.0.0, July 2023, URL: <https://github.com/oKermorgant/log2plot>.
- [59] F. C. P. Kevin M. Lynch, *Modern Robotics: Mechanics, Planning, and Control*, Cambridge, May 2017, 544 pp., ISBN: 1107156300.
- [60] H. Khalil, *Nonlinear systems*, Upper Saddle River, New Jersey: Prentice Hall, 2002, ISBN: 9780130673893.
- [61] B. Krauskopf et al., « A Survey of Methods for Computing (un)Stable Manifolds of Vector Fields », *in: Int. Journal of Bifurcation and Chaos* 15.03 (2005), pp. 763–791, DOI: 10.1142/S0218127405012533.
- [62] A. Kriegl and P. W. Michor, *The Convenient Setting of Global Analysis*, American Mathematical Society, 1997.
- [63] V. Kyrki, D. Kragic, and H. Christensen, « New shortest-path approaches to visual servoing », *in: IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS'04)*, Sendai, Japan, Sept. 2004, pp. 349–354.
- [64] J. P. LaSalle, « Some Extensions of Liapunov's Second Method », *in: IRE Transactions on Circuit Theory* 7.4 (1960), pp. 520–527, DOI: 10.1109/tct.1960.1086720.

-
- [65] J. Lee, *Introduction to topological manifolds*, vol. 202, Springer Science & Business Media, 2010.
- [66] J. M. Lee, *Introduction to Smooth Manifolds*, 2000.
- [67] E. Malis, « Improving vision-based control using efficient second-order minimization techniques », *in: IEEE Int. Conf. on Robotics and Automation (ICRA'04)*, New Orleans, LA, Apr. 2004, DOI: 10.1109/robot.2004.1308092.
- [68] E. Malis, F. Chaumette, and S. Boudet, « 2 1/2 D visual servoing », *in: IEEE Trans. on Robotics and Automation* 15.2 (Apr. 1999), pp. 238–250.
- [69] E. Malis and P. Rives, « Robustness of image-based visual servoing with respect to depth distribution errors », *in: IEEE International Conference on Robotics and Automation*, 2003, DOI: 10.1109/robot.2003.1241732.
- [70] E. Marchand, F. Spindler, and F. Chaumette, « ViSP for visual servoing: a generic software platform with a wide class of robot control skills », *in: IEEE Robot. Autom. Mag.* 12.4 (Dec. 2005).
- [71] B. Martin et al., « Certified Parallelotope Continuation for One-Manifolds », *in: SIAM Journal on Numerical Analysis* 51.6 (2013), pp. 3373–3401.
- [72] P. Martinet, J. Gallice, and D. Khadraoui, « Vision based control law using 3D visual features », *in: World Automation Congress*, Montpellier, France, May 1996, pp. 497–502.
- [73] H. Michel and P. Rives, *Singularities in the determination of the situation of a robot effector from the perspective view of three points*, tech. rep. 1850, Inria, Feb. 1993.
- [74] J. W. Milnor, *Morse theory*, 51, Princeton university press, 1963.
- [75] J. C. Moore, « A note on point-set mappings », *in: Papers in quantitative economics* 1 (1968), pp. 129–137.
- [76] J. A. Moreno, « Lyapunov Approach for Analysis and Design of Second Order Sliding Mode Algorithms », *in: Sliding Modes after the First Decade of the 21st Century: State of the Art*, ed. by L. Fridman, J. Moreno, and R. Iriarte, Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 113–149, ISBN: 978-3-642-22164-4, DOI: 10.1007/978-3-642-22164-4_4.

-
- [77] J. A. Moreno et al., « Multivariable Super-Twisting Algorithm for Systems with Uncertain Input Matrix and Perturbations », *in: IEEE Transactions on Automatic Control* (2021), pp. 1–1, DOI: 10.1109/tac.2021.3130880.
- [78] M. Morse, *The calculus of variations in the large*, vol. 18, American Mathematical Soc., 1934.
- [79] J. Nocedal and S. Wright, *Numerical optimization*, Springer Science & Business Media, 2006.
- [80] D. Oberkampf, D. Dementhon, and L. Davis, « Iterative pose estimation using coplanar feature points », *in: Computer Vision and Image Understanding* 63.3 (May 1996), pp. 495–511.
- [81] E. Outerelo and J. M. Ruiz, *Mapping Degree Theory*, American Mathematical Society, 2009.
- [82] J. Palis and W. de Melo, *Geometric Theory of Dynamical Systems*, Springer US, 1982, DOI: 10.1007/978-1-4612-5703-5.
- [83] B. N. Parlett, *The Symmetric Eigenvalue Problem*, Society for Industrial and Applied Mathematics, Jan. 1998, DOI: 10.1137/1.9781611971163.
- [84] B. Pascual-Escudero et al., « Complete Singularity Analysis for the Perspective-Four-Point Problem », *en, in: Int. Journal of Computer Vision* 129.4 (Apr. 2021), pp. 1217–1237, ISSN: 1573-1405, DOI: 10.1007/s11263-020-01420-0.
- [85] V. Pascucci et al., « Robust on-line computation of Reeb graphs », *in: ACM Transactions on Graphics* 26.3 (July 2007), p. 58, DOI: 10.1145/1276377.1276449.
- [86] K. B. Petersen and M. S. Pedersen, *The Matrix Cookbook*, Version 20121115, Nov. 2012, URL: <http://www2.compute.dtu.dk/pubdb/pubs/3274-full.html>.
- [87] M. Piras, *Application of interval analysis to the numerical investigation of properties of image-based visual servoing*, Master’s thesis, Aug. 2023.
- [88] A. Polyakov and A. Poznyak, « Reaching Time Estimation for “Super-Twisting” Second Order Sliding Mode Controller via Lyapunov Function Designing », *in: IEEE Transactions on Automatic Control* 54.8 (Aug. 2009), pp. 1951–1955, DOI: 10.1109/tac.2009.2023781.

-
- [89] V. H. P. Rodrigues and T. R. Oliveira, « Global adaptive HOSM differentiators via monitoring functions and hybrid state-norm observers for output feedback », *in: International Journal of Control* 91.9 (2018), pp. 2060–2072, DOI: 10.1080/00207179.2017.1392041.
- [90] C. Rucker, « Integrating Rotations Using Nonunit Quaternions », *in: IEEE Robotics and Automation Letters* 3.4 (Oct. 2018), pp. 2979–2986, DOI: 10.1109/lra.2018.2849557.
- [91] P. Schönemann, « A generalized solution of the orthogonal Procrustes problem », *in: Psychometrika* 31.1 (Mar. 1966), pp. 1–10, DOI: 10.1007/bf02289451.
- [92] F. Schramm et al., « Extended-2D visual servoing », *in: IEEE Int. Conf. on Robotics and Automation (ICRA'04)*, New Orleans, LA, Apr. 2004, pp. 267–273.
- [93] R. Seeber and M. Horn, « Stability proof for a well-established super-twisting parameter setting », *in: Automatica* 84 (2017), pp. 241–243, ISSN: 0005-1098, DOI: <https://doi.org/10.1016/j.automatica.2017.07.002>.
- [94] R. Seeber, M. Horn, and L. Fridman, « A Novel Method to Estimate the Reaching Time of the Super-Twisting Algorithm », *in: IEEE Transactions on Automatic Control* 63.12 (2018), pp. 4301–4308, DOI: 10.1109/TAC.2018.2812789.
- [95] D. S. Shafer, « Topological equivalence of gradient vectorfields », *in: Transactions of the American Mathematical Society* (1980).
- [96] Y. B. Shtessel et al., *Sliding mode control and observation*, Springer, New York, USA, 2014.
- [97] S. Smale, « Morse inequalities for a dynamical system », *in: Bulletin of American Mathematical Society* (1960).
- [98] M. Spivak, *Calculus on Manifolds*, Addison-Wesley Publishing Company, 1965.
- [99] E. Thompson, « Space resection: failure cases », *in: The Photogrammetric Record* 5.27 (Oct. 1966), pp. 201–204.
- [100] B. Thuilot et al., « Position based visual servoing: Keeping the object in the field of vision », *in: IEEE Int. Conf. on Robotics and Automation, ICRA'02*, Washington DC, May 2002, pp. 1624–1629.

-
- [101] B. Tibken and Y. Fan, « Computing the domain of attraction for polynomial systems via BMI optimization method », *in: 2006 American Control Conference*, IEEE, 2006, DOI: 10.1109/acc.2006.1655340.
- [102] G. Trombettoni et al., « Inner Regions and Interval Linearizations for Global Optimization. », *in: AAAI*, ed. by W. Burgard and D. Roth, AAAI Press, 2011.
- [103] W. Wilson, C. Hulls, and G. Bell, « Relative end-effector control using Cartesian position based visual servoing », *in: IEEE Trans. on Robotics and Automation* 12.5 (Oct. 1996), pp. 684–696.
- [104] V. Zubov, *Methods of A. M. Lyapunov and Their Application*, ed. by P. N. L. Leo F. Boron, 1964.

Titre : Stabilité et Régions d'Attraction de Systèmes d'Asservissement Visuel Basé Image

Mot clés : Asservissement visuel, analyse de la stabilité, régions d'attraction

Résumé : L'asservissement visuel est un domaine mature dont le formalisme est bien établi. L'un de ses principaux objectifs est le contrôle d'un robot à l'aide des informations visuelles provenant d'une caméra, qui est généralement montée sur l'effecteur du robot. Même si l'asservissement visuel a conduit à de nombreuses applications réussies, son analyse de stabilité reste une question théorique ouverte, en particulier pour l'asservissement visuel basé image. Dans cette thèse, nous avons exploré en profondeur les propriétés de stabilité de ces systèmes, à la fois en développant de nouveaux outils théoriques qui nous permettent d'étudier la stabilité des systèmes à *énergie décroissante* et en définissant une stratégie qui rend le calcul de l'ensemble complet des équilibres ré-

solvable dans le contexte de l'asservissement visuel. Grâce à notre nouvelle méthodologie, qui s'inspire des théories de Lyapunov et de Morse, nous avons pu caractériser rigoureusement l'interconnexion entre les équilibres de ces systèmes et proposer les régions d'attraction *indépendante des contrôleurs* des équilibres stables, ce qui a permis d'éclairer d'un jour nouveau le comportement de ces systèmes. De plus, les outils que nous avons développés nous ont permis de synthétiser un nouveau paradigme de contrôle, que nous appelons *meta-contrôleur*, qui fournit une amélioration substantielle de la taille de la région d'attraction de la pose souhaitée, mettant en évidence l'importance de la compréhension de la stabilité de ces systèmes, et l'importance des points selles pour son analyse.

Title: Stability and Regions of Attraction of Image-Based Visual Servoing Systems

Keywords: Visual servoing, stability analysis, regions of attraction

Abstract: Visual servoing is a mature area whose formalism is well established. One of its main aims is the control of a robot using the visual information coming from a camera, which is typically mounted as the end effector of the robot. Even if visual servoing has led to many successful applications, its stability analysis is still an open theoretical issue, in particular for image-based visual servoing. In this thesis, we deeply explored the stability properties of these systems, both by developing new theoretical tools that allow us to study the stability of *energy-decreasing* systems and by creating a strategy that makes the computation of the full set of equilibria tractable in the context of visual servoing. Thanks to our new methodol-

ogy, which is inspired by Lyapunov and Morse theories, we were able to rigorously characterize the interconnection between the equilibria of these systems and to propose the *controller-independent* regions of attraction of the stable ones, shining a new light on the behavior of these systems. Moreover, the tools that we developed allowed us to synthesize a new control paradigm, that we call *meta-controller*, which provides a substantial improvement in size for the desired pose's region of attraction, showcasing the importance of understanding the stability of these systems and the importance of saddles in their analysis.