



HAL
open science

Sequentially learning regions of attraction from data

Oumayma Khattabi, Matteo Tacchi, Sorin Olaru

► **To cite this version:**

Oumayma Khattabi, Matteo Tacchi, Sorin Olaru. Sequentially learning regions of attraction from data. MED 2025 - 33rd Mediterranean Conference on Control and Automation, Jun 2025, Tangier, Morocco. pp.589-594, <10.1109/MED64031.2025.11073289>. <hal-05059632>

HAL Id: hal-05059632

<https://hal.science/hal-05059632v1>

Submitted on 7 May 2025

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.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

Sequentially learning regions of attraction from data

Oumayma Khattabi, Matteo Tacchi-Bénard, Sorin Olaru

May 5, 2025

Abstract

The paper is dedicated to data-driven analysis of dynamical systems. It deals with certifying the basin of attraction of a stable equilibrium for an unknown dynamical system. It is supposed that point-wise evaluation of the right-hand side of the ordinary differential equation governing the system is available for a set of points in the state space. Technically, a Piecewise Affine Lyapunov function will be constructed iteratively using an optimisation-based technique for the effective validation of the certificates. As a main contribution, whenever those certificates are violated locally, a refinement of the domain and the associated tessellation is produced, thus leading to an improvement in the description of the domain of attraction.

Keywords : *Data-driven, Piecewise affine, Lyapunov function, Optimisation, Region of attraction, Iterative learning.*

1 Introduction

The ongoing rapid increase in data and computing capacity has led to significant advances in control theory, mainly with approaches exploiting data, gaining traction after important achievements in stabilising dynamical systems. Stability analysis follows two main currents: model-based and data-driven. While model-based methods (e.g. model-based predictive control (MPC), adaptive control, Lyapunov direct methods, etc.) have strong stability guarantees, data-driven control laws bypass the model identification step. They are directly building on to real system behaviour, a feature that promotes them as versatile alternatives in stabilising complex systems. However, the critical issue of control stability certification arises, hindering their applicability.

Traditionally, one of the key frameworks used in stability analysis is Lyapunov stability theory [1]. It is based on finding a function/functional related to the dynamical system, called the Lyapunov candidate, that decreases along trajectories of the system and converges to zero as the system converges to a given attractor (e.g. an equilibrium point). Constructing an adequate Lyapunov function using analytical methods for systems with simple models is moderately easy, but challenging for systems with complex and non-linear models. Some alternative solutions have been proposed; for example, in [2], the system function in question is represented by a piecewise affine (PWA) approximation for which a PWA Lyapunov candidate is computed using linear programming. In [3], on the other hand, a neural Lyapunov function is constructed to identify positively invariant sets using a PWA activation function. Neural Lyapunov functions are also used in [4] to certify regions of attraction of the controlled system. There are other methods than constructing Lyapunov functions, like the interval-based method in [5] and the interval Lyapunov equation test of [6], which are also used for stability analysis of complex systems.

Many works in the literature propose data-based methods to prove the stability of the equilibrium point of the system. A prominent method involves the use of deep neural networks. This can be found in [7] where a neural Lyapunov function is learnt along with the system's dynamics to ensure stability guarantees. Another example is [8], where a neural network-based control law is used and is constrained with a Lyapunov function learnt online. In [9], a neural Lyapunov candidate is

constructed while trying to avoid the curse of dimensionality related to neural networks. Another notable stability analysis approach is the continuous piecewise affine (CPA) method by [10], where data points are taken on vertices of a space partition and are used to construct the Lyapunov function. Beyond the approaches above, [11] focuses on Taylor polynomials and utilises piecewise polynomial approximations based on noisy data to study system dissipativity. Other recent research includes [12], where a safe stabilising control law is generated along with its Lyapunov function for discrete systems using a noisy dataset.

Existing data-driven stability analysis approaches provide various frameworks for studying system behaviour. This paper will examine one such approach, presented in [13], which seeks to find PWA Lyapunov functions for non-linear systems using data. The present contribution is organised as follows: After introducing the mathematical framework, the optimisation problem at the procedure’s core is recalled, and as a contribution, an iterative implementation scheme is proposed. The advantage of such a parsimonious treatment of the Lyapunov candidate construction is that it provides patchy Lyapunov functions [14] and thus allows for the certification of the domain of attraction in cases that are difficult to validate over the entirety of the region. We propose a technique that targets the uncertified area iteratively and identifies invariant sets based on the learnt Lyapunov functions. Finally, we present a successful resolution of the previous problematic case as a numerical application.

Notation: \mathbb{R}^+ is defined as the set $\{x \in \mathbb{R} | x \geq 0\}$. For a topological space \mathcal{X} , $\text{int}(\mathcal{X})$, $\overline{\mathcal{X}}$, and $\partial\mathcal{X}$ denote its interior, closure, and boundary, respectively. A polyhedron is the intersection of a finite number of half-spaces. A polyunion is a bounded, finite union of polyhedra (not necessarily convex). The convex hull of a polyunion is denoted as $\text{conv}(\cdot)$. For a set S , its cardinal is denoted as $|S|$. This notation is extended for a polyhedron P so that $|P|$ is the number of vertices in P . For $x \in \mathbb{R}^n$ and $r \in \mathbb{R}$, $\mathcal{B}(x, r)$ is used to denote a ball in \mathbb{R}^n of center x and radius r . The Euclidean norm is denoted as $\|\cdot\|$. Let f be a function from $\mathcal{X} \subset \mathbb{R}^n$ to \mathbb{R}^n , it is Lipschitz continuous, and M is its Lipschitz constant, if for all $x, y \in \mathcal{X}$, $\|f(x) - f(y)\| \leq M\|x - y\|$. The strict sub-level set of a function V for a certain level $\alpha \in \mathbb{R}$ is defined as $\mathbf{L}_\alpha^V := \{x \in \mathbb{R}^n | V(x) < \alpha\}$. An N_c -piece tessellation $\{Y_c\}_{1 \leq c \leq N_c}$ of a subspace $\mathcal{X} \subset \mathbb{R}^n$ is defined as $\bigcup_{c=1}^{N_c} Y_c = \overline{\mathcal{X}}$, such that $\forall c, c' \in \{1, \dots, N_c\}$, $c \neq c' \Rightarrow \text{int}(Y_c) \cap \text{int}(Y_{c'}) = \emptyset$, with Y_c convex polyunions which will be called simplexes. In the 2D case, the one-dimensional common facet between two simplices is called an edge, and a singular point of contact between edges is called a vertex. The set of vertices of the tessellation will be denoted $\{v_{i,c}\}_{\substack{1 \leq i \leq |Y_c| \\ 1 \leq c \leq N_c}}$.

2 Problem Statement

This section establishes the theoretical framework necessary for our stability analysis approach. First, we present the fundamental mathematical assumptions to lay the groundwork. Second, we reintroduce the optimisation problem given in [13] that constitutes the basis of our methodology. Last, we define the objectives and difficulties to tackle.

2.1 Framework

Our study addresses the challenge of analysing the domain of attraction of the system with unknown dynamics :

$$\dot{x} = f(x) \tag{1}$$

In order to set up a framework that allows a feasible solution, a series of hypotheses have to be imposed with respect to the family of dynamics and the region over which the analysis will be carried out.

Assumption 1. The system under study has a minimal state space realisation. For such a minimal state space realisation, the following hold:

- the dimension n of the state vector is known;
- the right-hand side $f(\cdot)$ of (1) is accessible for evaluation either via measurements or through simulation;
- the function $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is Lipschitz continuous, and an upper bound M for its Lipschitz constant is known a priori.

Analysing dynamical systems can reveal complex behaviours, including multiple equilibria and chaotic trajectories. To maintain focus within our study's scope, we introduce additional constraints defining the solution space and the qualitative/quantitative structure of the solutions to (1).

Assumption 2. The analysis will be carried on on a bounded set $\mathcal{X} \subset \mathbb{R}^n$ of the state space. It is assumed that there is one locally asymptotically stable equilibrium $\bar{x} = 0 \in \text{int}(\mathcal{X})$, and no other equilibrium point exists in \mathcal{X} .

Based on Assumption 1, we can have access to a dataset $\mathcal{D} := \{(x_d, f_d)\}_{1 \leq d \leq N_d}$, with $f_d = f(x_d)$ and N_d the number of data points, and we can choose the position of the state samples x_d when generating data.

The aim is to be able to certify the stability of system (1) in \mathcal{X} based on the dataset \mathcal{D} and the Lipschitz constant M .

2.2 Data-driven Lyapunov stability analysis

This work is based on the method presented by Tacchi et al. in [13], in which we identify regions of attraction by constructing certified and robust Lyapunov functions consistent with the available data. The following Assumptions provide two additional technical conditions to avoid singularities in the numerical implementation:

Assumption 3. The following information is available with respect to (1):

- a neighbourhood \mathcal{A} of the equilibrium point is known to be attracted to $\bar{x} = 0$;
- a fixed dataset \mathcal{D} of the system is available ;
- $\bar{\mathcal{X}}$ and \mathcal{A} are polyunions ;

Let the region $\bar{\mathcal{X}} \setminus \bar{\mathcal{A}}$ be subdivided into an N_c -piece tessellation $\{Y_c\}_{1 \leq c \leq N_c}$, with N_c the number of cells Y_c . We set each cell Y_c to be a convex polyunion with vertices $\{v_{1,c}, \dots, v_{|Y_c|,c}\}$. To each cell Y_c , an affine Lyapunov function $V_c(x) := g_c^\top x + b_c$ is assigned. These functions are continuously merged together and constitute one global PWA Lyapunov candidate V on \mathcal{X} defined as follows :

$$\forall c \in \{1, \dots, N_c\}, \forall x \in Y_c \subset \bar{\mathcal{X}} \setminus \bar{\mathcal{A}}, \quad V(x) = g_c^\top x + b_c \quad (2)$$

Based on these elements, and according to [13, Proposition 6], if the upper bound on $\nabla V^\top f$ based on data is negative at all points in $\bar{\mathcal{X}} \setminus \bar{\mathcal{A}}$ where ∇V exists, then for all $\alpha \in \mathbb{R}$ such that $\mathbf{L}_\alpha^V \subset \mathcal{X}$, \mathbf{L}_α^V is include in the region of attraction.

The following optimisation problem is a formulation of the conditions to find a PWA Lyapunov function V :

$$\mathbf{s}_{\alpha,\mu} := \min_{\{\gamma_{d,c}, g_c, b_c, s_{i,c}\}} \sum_{c=1}^{N_c} \sum_{i=1}^{|Y_c|} s_{i,c} \quad (3a)$$

$$\forall c, c' \in \{1, \dots, N_c\}, i \in \{1, \dots, |Y_c|\}, s_{i,c} \geq -\mu \quad (3b)$$

$$v_{i,c} \in Y_{c'} \Rightarrow (g_c - g_{c'})^\top v_{i,c} = b_{c'} - b_c \quad (3c)$$

$$\sum_{d=1}^{N_d} \gamma_{d,c} = g_c \quad (3d)$$

$$\sum_{d=1}^{N_d} \gamma_{d,c}^\top f_d + \|\gamma_{d,c}\| M \|v_{i,c} - x_d\| \leq s_{i,c} \quad (3e)$$

with $\{s_{i,c}\}_{1 \leq i \leq |Y_c|, 1 \leq c \leq N_c}$ some slack variables to minimise (one per vertex of each cell of the tessellation), $\{\gamma_{d,c}\}_{1 \leq d \leq N_d, 1 \leq c \leq N_c}$ data dependent variables that tune $\{g_c\}_{1 \leq c \leq N_c}$ using constraints (3d), and $\mu \in \mathbb{R}^+$ defining the lower bound of the slack variables. The constraints (3c) represent the continuity of the Lyapunov function V between neighbouring cells Y_c and $Y_{c'}$. The negativity of $\nabla V^\top f$ is ensured by constraint (3e). Once the function is obtained, a level α should be found such that $\mathbf{L}_\alpha^V \subset \mathcal{X}$ with $\mathcal{A} \subset \mathbf{L}_\alpha^V$, proving that \mathbf{L}_α^V is attracted to the equilibrium. The learned function V is a certified Lyapunov function if all slack variables $\{s_{i,c}\}_{1 \leq i \leq |Y_c|, 1 \leq c \leq N_c}$ are negative. The negativity of $\nabla V^\top f$ is thus guaranteed, making the resulting V a valid Lyapunov function.

To avoid the problem of infeasibility, [13] provides a necessary condition to be validated by the dataset :

$$\{v_{i,c}\}_{1 \leq i \leq |Y_c|, 1 \leq c \leq N_c} \subset \bigcup_{d=1}^{N_d} \mathcal{B} \left(x_d, r_d = \frac{\|f_d\|}{M} \right) \quad (4)$$

This implies that data should provide at least some minimal information relevant to each vertex of the tessellation.

2.3 Advantages, challenges and objectives

This method is a promising alternative to the state of the art, where the common practice is to use semi-definite programming (SDP) [11, 12]. In contrast, the optimisation problem (3) is based on second-order cone programming (SOCP) and linear constraints, which makes it more scalable to higher dimensions. Also, unlike neural networks that train on big datasets, it does not need ample datasets to give satisfying results. Moreover, there is no binding relation between the position of the data points and the tessellation except for the learnability condition (4), which is quite flexible. This gives more freedom to adjust the dataset and vertices depending on the system at hand.

It should be mentioned that optimisation problem (3) gives satisfying results when applied to various nonlinear systems with nonconvex regions of attraction, like the Van Der Pol oscillator (see numerical results in [13]). In some cases, constraints (3b) are not saturated, but all slack variables are negative, which is enough to certify attractivity. However, there are exceptional cases where numerical computations provide (according to the chosen tessellation) a Lyapunov function certified in only part of the set considered initially. An example of one such case is the damped pendulum with the following state space representation:

$$\dot{x} = \begin{pmatrix} \dot{\theta} \\ \dot{\omega} \end{pmatrix} = \begin{pmatrix} \omega \\ -\sin(\theta) - 2\omega \end{pmatrix} \quad (5)$$

For a randomly chosen Delaunay triangulation and dataset as displayed in Figure 1a, the stability is, on average, not certified by the algorithms of [13] in the diagonal central part of the region (see figure 1b), even though it is attracted to the equilibrium. The expected result was a Lyapunov function covering the entirety of the tessellation. However, the optimisation problem could not provide negative slack variables and thus valid affine functions in some of the simplexes. These simplexes get eliminated as they don't validate the Lyapunov function criteria, which leaves a considerable area of the initial set uncovered and thus uncertified by the PWA Lyapunov candidate. This result is hypothesised to be caused by the drastic variation of the values of $\{f_d\}_{1 \leq d \leq N_d}$ (represented by orange arrows in 1a), where data points with larger image values (big arrows in the corners) exert significant influence on the optimisation process but fail to sufficiently contribute information on the images of the vertices in areas where f_d takes on smaller absolute values. Consequently, these areas with lower function absolute values are under-represented and inaccurately characterised, eventually affecting the overall accuracy of the optimisation.

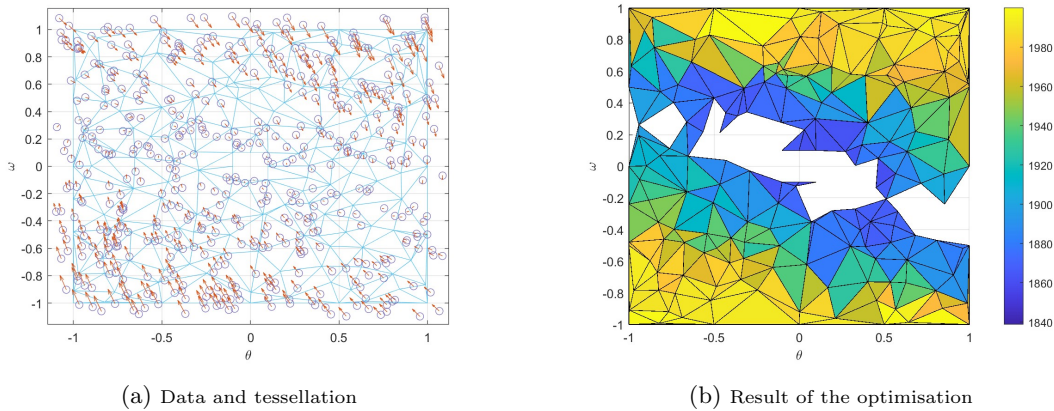


Figure 1: Example case of the simple pendulum

According to [13, Appendix B], a dataset and a tessellation providing enough covering of the set $\mathcal{X} \setminus \mathcal{A}$ are bound to bring forth a certified Lyapunov function. However, there is no sufficient condition for the density and distribution of the data and vertices to allow the algorithm's convergence to optimality. Moreover, the complexity of the computations explodes as the number of constraints is of the order of $\mathcal{O}((3|Y_c| + 1) \cdot N_c)$ and the number of optimisation variables is around $\mathcal{O}((|Y_c| + 2N_d + 3) \cdot N_c)$, which makes it difficult to boundlessly increase N_d and N_c until optimality is reached.

The objective of this work is to solve the non-certifiability of the Lyapunov function in parts of the region and to overcome the complexity issue to ensure that solving this optimisation problem can effectively establish the stability of the equilibrium point of the system in \mathcal{X} , using the example of the pendulum of system (5) as validation.

3 Proposed Approach

To overcome the challenges encountered when implementing this stability analysis method and demonstrate the practical applicability of the mathematical theory, we propose in this section an iterative approach using the optimisation problem, where we iteratively adapt the data and tessellation to deploy to get a Lyapunov function certifying the largest region of attraction to the equilibrium point in the set \mathcal{X} .

3.1 The principle of the iterative construction

Since one of the core issues lies in the uncertainty regarding the existence of a PWA function solution for the optimisation problem with all slacks negative, a first sanity check is needed to ascertain that a PWA Lyapunov candidate exists. Note that this sanity check uses the full knowledge of the system model (5), while usually, the dynamics f are unknown. However, our iterative construction still applies in the unknown setting, provided one can refine their tessellation and data distribution.

For this, we consider a quadratic Lyapunov function V_Q of the system (5):

$$\forall x \in \mathbb{R}^2 \text{ s.t. } \|x\|_\infty \leq 1, \quad V_Q(x) = \frac{1}{2}x^T Q x \quad (6)$$

with $Q = \begin{pmatrix} 10 & 3 \\ 3 & 4 \end{pmatrix}$.

Next, a PWA approximation \hat{V}_Q of V_Q is computed. If \hat{V}_Q is a Lyapunov candidate, then we have proven the existence of a PWA Lyapunov function and a tessellation/subdivision $\{Y_c\}_{1 \leq c \leq N_c}$ corresponding to the affine pieces of \hat{V}_Q .

Since we know the model and a PWA Lyapunov function here, we can use the corresponding subdivision as our initial tessellation for faster convergence. In the general case, when the dynamics are unknown, using a Delaunay triangulation also works [15], but it only takes more iterations to converge. Next, we generate data points that offer a covering of the tessellation so that learnability condition (4) is met. Then, we solve optimisation problem (3) for the constructed inputs, obtain a Lyapunov candidate V , and look for non-negative slack variables. If none exist, then the largest α with $\mathcal{A} \subset \mathbf{L}_\alpha^V \subset \mathcal{X}$ defines a region attracted to 0 along system trajectories. Else, we identify the uncertified points

$$\mathcal{N} = \{(i, c) \in \mathbb{N}^2 \mid 1 \leq c \leq N_c, 1 \leq i \leq |Y_c|, s_{i,c} \geq 0\}$$

from which we deduce a convex uncertified region:

$$\mathcal{C} = \text{conv}(\mathcal{A} \cup \{v_{i,c}\}_{(i,c) \in \mathcal{N}}). \quad (7)$$

By design of constraint (3e) (see [13, Appendix A] for details), we can certify that $\forall x \in \mathcal{X} \setminus \mathcal{C}$, $\nabla V(x)^\top f(x) < 0$. The problem is that, contrary to the setting in [13, Appendix A], the case $\mathcal{C} \not\subseteq \mathcal{A}$ can occur, preventing us from concluding anything of significance on whether \mathcal{C} is part of the region of attraction to the equilibrium. In that case, one can study the stability in \mathcal{C} by iterating as follows:

- suppose a first iteration k failed to compute V_k , α_k and \mathcal{C}_k based on \mathcal{X}_k , \mathcal{A}_k and \mathcal{D}_k such that $\mathcal{C}_k \not\subseteq \mathcal{A}$. Define $\mathcal{X}_{k+1} = \beta_{k+1} \cdot \mathcal{C}_k$, where $\beta_{k+1} \geq 1$ is the highest upscaling parameter with $\mathcal{X}_{k+1} \subset \mathbf{L}_{\alpha_k}^{V_k}$ still verified; if no such $\beta_{k+1} \geq 1$ exists, it means that the uncertified area intersects or is outside of the level set limit where $V_k = \alpha_k$, meaning this particular iteration failed and should be restarted with a finer tessellation and more data points until a better and valid result is generated.
- Introduce a smaller $\mathcal{A}_{k+1} \subset \mathcal{A}_k \subset \mathcal{C}_k$ ($\mathcal{A}_0 = \mathcal{A}$ ensures all such sets are attracted to the equilibrium, and non-negative slacks are often in the neighbourhood of the exclusion zone \mathcal{A}_k , so taking a smaller \mathcal{A}_{k+1} can validate the condition \mathcal{C}_{k+1} in \mathcal{A}).
- Compute a tessellation of $\mathcal{P} = \mathcal{X}_{k+1} \setminus \mathcal{A}_{k+1}$ by adding vertices to make it denser, and refine the dataset \mathcal{D}_k into some \mathcal{D}_{k+1} accordingly by removing some data points and adding others while maintaining learnability condition (4).
- Solve optimisation problem (3) for the new input to obtain a new V_{k+1} . Find α_{k+1} such that $\mathcal{C}_k \subset \mathbf{L}_{\alpha_{k+1}}^{V_{k+1}} \subset \mathcal{X}_{k+1}$; if no α_{k+1} such that $\mathcal{C}_k \subset \mathbf{L}_{\alpha_{k+1}}^{V_{k+1}}$ exists, the new Lyapunov candidate fails to cover all uncertified points and the iteration fails. This iteration needs to be done again in this case with refined tessellation and richer data. Otherwise, define \mathcal{C}_{k+1} and proceed.

- If $\mathcal{C}_k \subset \mathcal{A}$ or $k = k_{\max}$ (max number of iterations reached), stop iterating; else, start again from the first step with $k \leftarrow k + 1$.

The iterative method consists of repeating all the steps, from choosing the parameters, data and tessellation to testing if the uncertified area is contained in \mathcal{A} , until this last condition is verified. It works following Algorithm 1.

Algorithm 1:

Input: $\mathcal{X}, \mathcal{A}, \mu, M$

Output: $\{\mathbf{L}_{\alpha_k}^{V_k}\}$

$\mathcal{P} \leftarrow \mathcal{X} \setminus \mathcal{A}$

while number of iterations $k < k_{\max}$ **do**

 Generate tessellation $\{Y_c\}_{1 \leq c \leq N_{c_k}}$ with $\mathcal{P} = \bigcup_{c=1}^{N_{c_k}} Y_c$

 Generate dataset \mathcal{D}_k verifying condition (4)

 Run optimisation problem (3) and obtain V_k

if all slack variables are negative **then**

 | Break

else

 Identify the non-certified polyunion \mathcal{C}_k

 Identify the corresponding level set $\mathbf{L}_{\alpha_k}^{V_k}$ of V_k

if $\mathcal{C}_{k-1} \not\subseteq \mathbf{L}_{\alpha_k}^{V_k}$ **then**

 | Iteration failed, restart with a more refined tessellation and dataset

else if $\mathcal{C}_k \subseteq \mathcal{A}_k$ **then**

 | Break

else if β_{k+1} exists such that $\beta_{k+1} \cdot \mathcal{C}_k \subset \mathbf{L}_{\alpha_k}^{V_k}$ **then**

 | Choose $\mathcal{A}_{k+1} \subset \mathcal{A}_k$

 | Accept β_{k+1} and $\mathcal{X}_{k+1} = \beta_{k+1} \cdot \mathcal{C}_k$

 | $\mathcal{P} \leftarrow \mathcal{X}_{k+1} \setminus \mathcal{A}_{k+1}$

end

end

Algorithm 1 comes with a certification result, which we state in Theorem 1.

Theorem 1. Let $\mathcal{X} = \mathcal{X}_0 \supset \mathcal{X}_1 \supset \dots \supset \mathcal{X}_K$, $\mathcal{A} = \mathcal{A}_0 \supset \mathcal{A}_1 \supset \dots \supset \mathcal{A}_K$, $\mathcal{D} = \mathcal{D}_0, \dots, \mathcal{D}_K$ defining ambient sets, prior attracted sets and datasets of instances of problem (3). Let V_0, \dots, V_K be the patchy PWA Lyapunov candidates associated to feasible solutions of the corresponding problem instances, with uncertified regions $\mathcal{C}_0, \dots, \mathcal{C}_K$ and levels $\alpha_0, \dots, \alpha_K$ (in particular, $\forall k \in \{0, \dots, K\}$, $\mathbf{L}_{\alpha_k}^{V_k} \subset \mathcal{X}_k$). Suppose that:

- $\forall k \in \{1, \dots, K\}$, $\mathcal{C}_{k-1} \subset \mathbf{L}_{\alpha_k}^{V_k} \subset \mathcal{X}_k \subset \mathbf{L}_{\alpha_{k-1}}^{V_{k-1}}$
- $\mathcal{C}_K = \mathcal{A}_K$ (no non-negative slack variable) or $\mathcal{C}_K \subset \mathcal{A}_0$

Then, $\mathbf{L}_{\alpha_0}^{V_0}$ is attracted to 0 along the unknown dynamics, i.e.

$$V_0(x(0)) \leq \alpha_0 \implies x(t) \xrightarrow[t \rightarrow \infty]{} 0. \quad (8)$$

Proof. Suppose that $K = 0$. Then, $\mathcal{C}_0 = \mathcal{A}_0$ i.e. all slack variables in problem (3) are negative, and [13, Theorem 8] ensures that (8) holds.

Suppose that $V_k(x(0)) \leq \alpha_k \implies \lim_{t \rightarrow \infty} x(t) = 0$. Then, since $\mathcal{C}_{k-1} \subset \mathbf{L}_{\alpha_k}^{V_k} \subset \mathcal{X}_{k-1}$ and V_{k-1} is a feasible solution for problem (3), $\nabla V_{k-1}(x)^\top f(x) < 0$ is certified for all $x \in \mathcal{X}_{k-1} \setminus \mathbf{L}_{\alpha_k}^{V_k}$ and according to [13, Proposition 6], it holds $V_{k-1}(x(0)) \leq \alpha_{k-1} \implies \lim_{t \rightarrow \infty} x(t) = 0$.

By induction, since $\mathbf{L}_{\alpha_0}^{V_0} \supset \mathbf{L}_{\alpha_1}^{V_1} \supset \dots \supset \mathbf{L}_{\alpha_K}^{V_K}$, it is therefore enough to prove that

$$V_K(x(0)) \leq \alpha_K \implies \lim_{t \rightarrow \infty} x(t) = 0,$$

which is done by noticing that, since $\mathcal{C}_K \subset \mathcal{A}_0$ (because $\mathcal{C}_K \subseteq \mathcal{A}_k \subset \mathcal{A}_0$), $\nabla V_K(x)^\top f(x) < 0$ is certified for all $x \in \mathcal{X}_K \cap (\mathcal{X}_0 \setminus \mathcal{A}_0)$ and by reusing [13, Proposition 6]. \square

3.2 Implementation and practical challenges

In [13], the last algorithm proposed (Algorithm 3) bears some similarity to our proposed Algorithm 1 as it is also an iterative approach, but they are fundamentally different. In [13, Algorithm 3], the iterations start from the smallest region to the largest to build one continuous Lyapunov function covering the union of the areas. It also stops if it cannot generate a certified Lyapunov function in the region. Our algorithm, on the other hand, starts from the global region and certifies a region of attraction towards an inner region. If the inner region is not included in \mathcal{A} , it repeats the same step while covering smaller and smaller areas in each iteration. The goal is not to build a global Lyapunov function, but to prove the region’s attraction iteratively by constructing patchy Lyapunov functions in the region such that each function defines an invariant set based on its sublevel set.

To provide greater flexibility in the optimisation problem and ensure the validation of continuity constraints, all the tessellations are taken so that the cells are triangular. As for the dataset, considering that the region that proved difficult for the algorithm has data that can’t be informative enough for all the vertices, we include data from neighbouring areas outside of the tessellation until condition (4) is satisfied. Some data points are also taken on the vertices for even better coverage.

4 Numerical Results

In this part, we present the implementation of the iterative method of the optimisation problem described in Algorithm 1 on the damped pendulum with the mathematical model (5). The tessellation used is ensured to contain a PWA Lyapunov function. We seek to retrieve the attractivity of the origin as the equilibrium in the chosen region using Lyapunov candidates computed from data measurements, ignoring the full model (5). It should be mentioned that even if the next simulation results shown are informed of the model solely in establishing the tessellation, a simulation that respects the assumption of the unavailability of the model at all times is also successful, albeit it needs more iterations. The results are shown for comparison in Figure 7.

For the first implementation, we use the parameters described in Table 1 such that the initial set \mathcal{X} is bounded by $l_{\mathcal{X}}$ and the neighbourhood \mathcal{A} of the equilibrium point is bounded by $l_{\mathcal{A}}$.

Table 1: Initial optimization parameters

| Parameter | Value |
|-------------------|-----------------------|
| $l_{\mathcal{X}}$ | $[-1, 1; -1, 1]$ |
| $l_{\mathcal{A}}$ | $0.1.l_{\mathcal{X}}$ |
| μ | 5 |
| M | 2.5 |

The algorithm was coded using YALMIP and MPT3 toolboxes [16, 17], was run on a laptop with 32GB memory and AMD Ryzen 7 7735U, and the solver used was MOSEK. The problem was solved in five iterations. Figures 2 to 6 show the results of each iteration.

Each figure represents the results of an iteration k . The left sub-figure represents the Lyapunov function V_k solution of the optimisation problem. The right sub-figure represents the following:

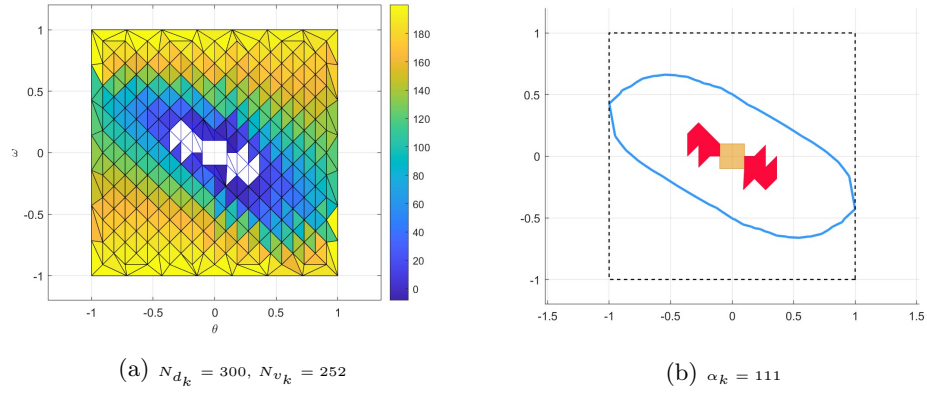


Figure 2: $k = 0$

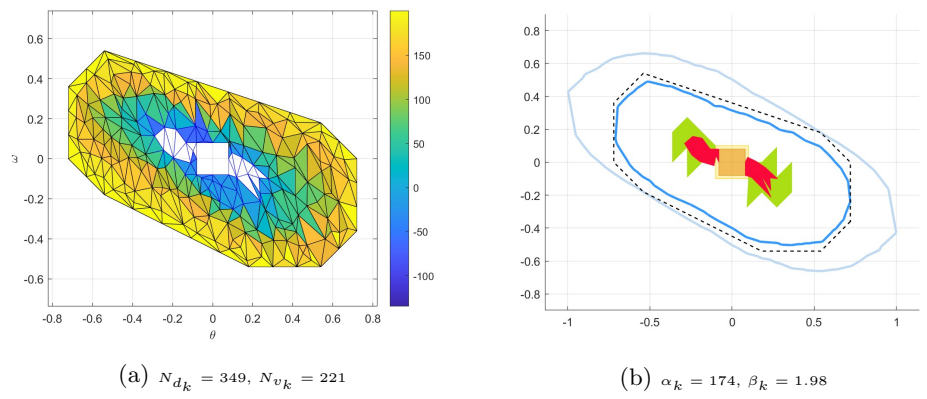


Figure 3: $k = 1$

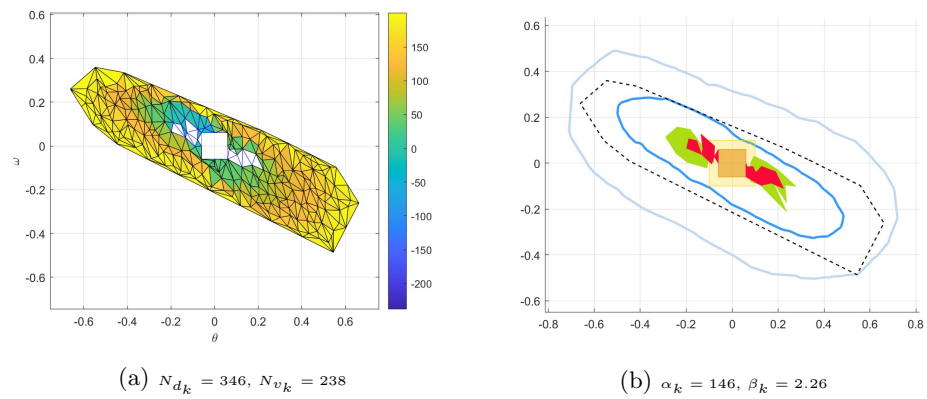


Figure 4: $k = 2$

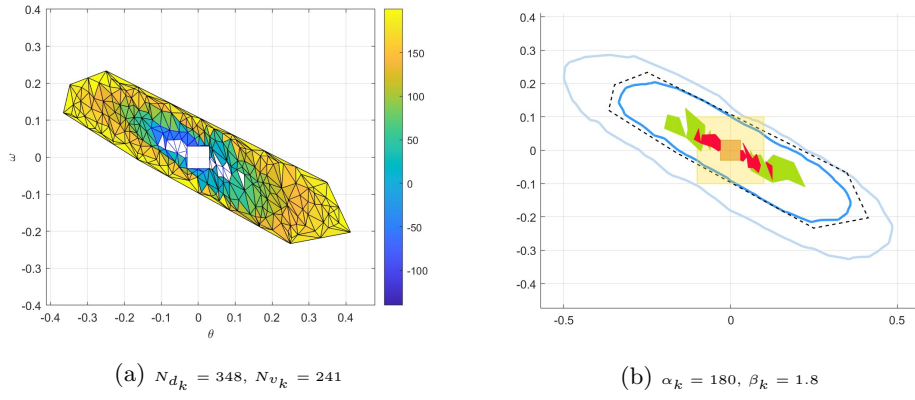


Figure 5: $k = 3$

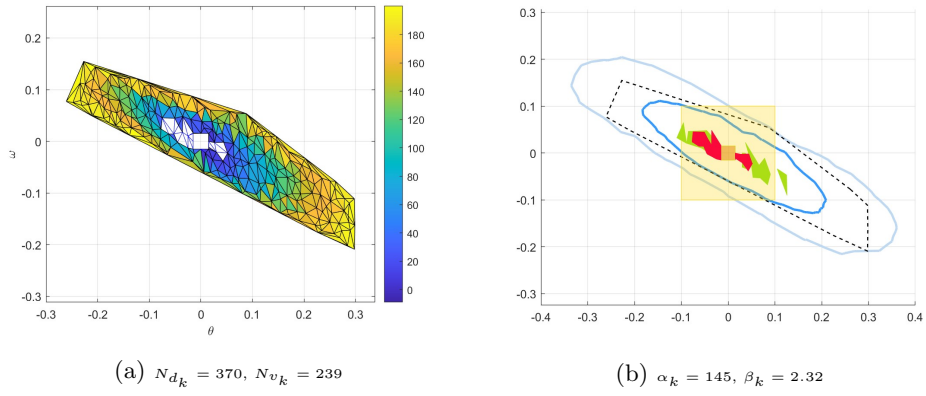


Figure 6: $k = 4$

- in a dashed line the borders of \mathcal{X}_k ;
- in a blue line the patchy region of attraction (delimited by the border of the level set $\mathbf{L}_{\alpha_k}^{V_k}$);
- in a faded blue line the border of the previous level set $\mathbf{L}_{\alpha_{k-1}}^{V_{k-1}}$;
- in red, the area \mathcal{C}_k that still needs to be certified;
- in green the previous uncertified area \mathcal{C}_{k-1} ;
- in orange the set \mathcal{A}_k ;
- and in transparent yellow the set $\mathcal{A} = \mathcal{A}_0$.

One can notice that the new level set $\mathbf{L}_{\alpha_k}^{V_k}$ inside \mathcal{X}_k in each iteration k includes the region \mathcal{C}_{k-1} not certified by the Lyapunov candidate in the previous step, which is in line with the first condition of Theorem 1: $\mathcal{C}_{k-1} \subset \mathbf{L}_{\alpha_k}^{V_k} \subset \mathcal{X}_k \subset \mathbf{L}_{\alpha_{k-1}}^{V_{k-1}}$. One can also see that \mathcal{C}_k covers less area than \mathcal{C}_{k-1} , hinting at algorithm convergence. \mathcal{C}_4 in particular is small enough to validate the stopping criterion, which is the inclusion in \mathcal{A} , in line with the validation of the second condition of Theorem 1: $\mathcal{C}_K \subset \mathcal{A} = \mathcal{A}_0$. The validation of the two conditions implies that the level set $\mathbf{L}_{\alpha_0}^{V_0}$ is a region of attraction.

Figure 7 shows the progression of the level sets from figures 2 to 6 over the iterations (from green to blue). All trajectories starting in an invariant set converge to the subset inside it. As for the last invariant set, the trajectories converge to \mathcal{A} and then to the equilibrium point. In grey dotted lines, we can see the level sets obtained from a completely data-driven simulation. It converged to \mathcal{A} successfully, albeit in one more iteration compared to the informed simulation above. This proves that it is possible to certify the attraction of the equilibrium point in the region without relying on the model, and that the choice of tessellation strongly influences the speed of convergence.

5 Conclusion

The paper concentrates on the data-driven stability certification. It employs an optimisation-based technique to build a PWA Lyapunov function and, as a contribution, proves that iterative treatment of the regions to be analysed can provide an increased complexity of the tessellations in the areas where the data are informative on the dynamical evolution. Practically, the resulting certificates of convergence rely on patchy Lyapunov functions and the construction of regions of attraction is ordered in the sense of inclusion. Several directions remain open for future research. Considering data's influence on the results, it would be interesting to consider some system trajectories as input instead of just independent data points. An essential next step is considering discrete and controlled systems to broaden the approach's applicability.

References

- [1] A. M. Lyapunov, "The general problem of the stability of motion," *International Journal of Control*, vol. 55, pp. 531–534, Mar 1992.
- [2] P. Julian, J. Guivant, and A. Desages, "A parametrization of piecewise linear Lyapunov functions via linear programming," *International Journal of Control*, vol. 72, pp. 702–715, Jan 1999.
- [3] D. Kim and H. J. Kim, "Estimation of constraint admissible invariant set with neural Lyapunov function," Sept. 2024. arXiv:2409.19881 [eess].




Figure 7: Graphical description of the relationship between the level sets' boundaries obtained iteratively. The progression in terms of set inclusions illustrates the convergence of the proposed algorithm.

- [4] J. Wang and M. Fazlyab, “Actor-critic physics-informed neural Lyapunov control,” *IEEE Control Systems Letters*, vol. 8, pp. 1751–1756, 2024. arXiv:2403.08448 [cs].
- [5] T. Le Mézo, L. Jaulin, and B. Zerr, “An interval approach to compute invariant sets,” *IEEE Transactions on Automatic Control*, vol. 62, pp. 4236–4242, Aug. 2017.
- [6] A. Goldsztejn and G. Chabert, “Estimating the robust domain of attraction for non-smooth systems using an interval Lyapunov equation,” *Automatica*, vol. 100, pp. 371–377, Feb. 2019.
- [7] J. Z. Kolter and G. Manek, “Learning stable deep dynamics models,” *Advances in neural information processing systems*, vol. 32, 2019.
- [8] Y. Min, S. M. Richards, and N. Azizan, “Data-driven control with inherent lyapunov stability,” in *2023 62nd IEEE Conference on Decision and Control (CDC)*, pp. 6032–6037, IEEE, 2023.
- [9] L. Grüne, “Computing Lyapunov functions using deep neural networks,” *arXiv preprint arXiv:2005.08965*, 2020.
- [10] S. F. Hafstein, C. M. Kellett, and H. Li, “Computing continuous and piecewise affine Lyapunov functions for nonlinear systems,” *Journal of Computational Dynamics*, vol. 2, no. 2, pp. 227–246, 2016.
- [11] T. Martin and F. Allgöwer, “Data-driven system analysis of nonlinear systems using polynomial approximation,” *IEEE Transactions on Automatic Control*, vol. 69, pp. 4261–4274, Jul 2024. Conference Name: IEEE Transactions on Automatic Control.
- [12] J. Zheng, J. Miller, and M. Sznaier, “Data-Driven Safe Control of Discrete-Time Non-Linear Systems,” *IEEE Control Systems Letters*, vol. 8, pp. 1553–1558, 2024. Conference Name: IEEE Control Systems Letters.
- [13] M. Tacchi, Y. Lian, and C. Jones, “Robustly learning regions of attraction from fixed data,” *IEEE Transactions on Automatic Control*, vol. 70, no. 3, pp. 1576–1591, 2025.
- [14] F. Ancona and A. Bressan, “Stabilization by patchy feedbacks and robustness properties,” in *Optimal control, stabilization and nonsmooth analysis*, pp. 185–199, Springer, 2004.
- [15] B. Delaunay, “Sur la sphère vide. À la mémoire de Georges Voronoï,” *Proceedings of the Russian Academy of Sciences. Mathematical Series*, no. 6, pp. 793–800, 1934.
- [16] J. Löfberg, “Yalmip : A toolbox for modeling and optimization in matlab,” in *In Proceedings of the CACSD Conference*, (Taipei, Taiwan), 2004.
- [17] M. Herceg, M. Kvasnica, C. Jones, and M. Morari, “Multi-parametric toolbox 3.0,” in *Proc. of the European Control Conference*, (Zürich, Switzerland), pp. 502–510, Jul 17–19 2013. <http://control.ee.ethz.ch/~mpt>.