

Counterexample-guided computation of polyhedral Lyapunov functions for piecewise linear systems [★]

Guillaume O. Berger ^a, Sriram Sankaranarayanan ^a,

^a*Department of Computer Science, University of Colorado, Boulder, USA*

Abstract

This paper presents a counterexample-guided iterative algorithm to compute convex, piecewise linear (polyhedral) Lyapunov functions for continuous-time piecewise linear systems. Polyhedral Lyapunov functions provide an alternative to commonly used polynomial Lyapunov functions. Our approach first characterizes intrinsic properties of a polyhedral Lyapunov function including its “eccentricity” and “robustness” to perturbations. We then derive an algorithm that either computes a polyhedral Lyapunov function proving that the system is asymptotically stable, or concludes that no polyhedral Lyapunov function exists whose eccentricity and robustness parameters satisfy some user-provided limits. Significantly, our approach places no a-priori bound on the number of linear pieces that make up the desired polyhedral Lyapunov function. The algorithm alternates between a learning step and a verification step, always maintaining a finite set of witness states. The learning step solves a linear program to compute a candidate Lyapunov function compatible with a finite set of witness states. In the verification step, our approach verifies whether the candidate Lyapunov function is a valid Lyapunov function for the system. If verification fails, we obtain a new witness. We prove a theoretical bound on the maximum number of iterations needed by our algorithm. We demonstrate the applicability of the algorithm on numerical examples.

Key words: Stability analysis, Lyapunov functions, Counterexample-guided methods, Linear optimization, Hybrid systems

1 Introduction

We study the problem of synthesizing Lyapunov functions for continuous-time, possibly uncertain, piecewise linear systems. These systems appear naturally in a wide range of applications (e.g., electrical circuits, mechanical systems with impact) or as approximations of more complex dynamical systems (Christoffersen, 2007; Xu and Xie, 2014). The existence of a Lyapunov function guarantees global convergence of the system toward the origin. However, finding such a Lyapunov function can be very challenging (Blondel and Tsitsiklis, 1999).

In this paper, we focus on Lyapunov functions that are convex and piecewise linear, also called *polyhedral* functions. Polyhedral functions are interesting because they

can approximate convex, positively homogeneous functions arbitrarily well. In fact, for a large class of hybrid systems, including switched linear systems, there exist converse results showing that if the system is asymptotically stable, then a polyhedral Lyapunov function exists (Sun and Ge, 2011). However, the computation of a polyhedral Lyapunov function with a given number of linear pieces is known to be difficult, even for linear systems (Blanchini and Miani, 2015). Furthermore, there are in general no a-priori bounds on the number of pieces that the function must have in order to be a Lyapunov function for a given class of systems (Ahmadi and Jungers, 2016).

In this paper, we present an iterative approach that searches for polyhedral Lyapunov functions for piecewise linear systems. At each iteration, the number of linear pieces for the desired polyhedral Lyapunov function is increased and the conditions that must be satisfied by these linear pieces are checked in an efficient manner (Section 6). The latter is achieved by providing a convex approximation of the Lyapunov constraints, obtained by enforcing the Lyapunov constraints only at a finite set of points in the state space. By associating a linear piece to each of these points, the computation of

[★] This research was funded in part by the Belgian-American Education Foundation (BAEF) and the US National Science Foundation (NSF) under award numbers 1836900 and 1932189.

Email addresses: guillaume.berger@colorado.edu (Guillaume O. Berger), sriram.sankaranarayanan@colorado.edu (Sriram Sankaranarayanan).

the polyhedral function can be formulated as a convex optimization program. The process terminates successfully, yielding a polyhedral Lyapunov function, or fails to find a Lyapunov function which lies within a class described by a single parameter, called the *eccentricity* (analogous to the eccentricity of an ellipsoid). Due to the convex approximation of the Lyapunov constraints, failure does not necessarily imply that no polyhedral Lyapunov function in this class exists for the system. However, one can conclude that the system does not admit a polyhedral Lyapunov function in this class that satisfies a robustness property specified by two parameters, called the *time step* and the *contraction rate* (Section 5). We further show that a failure also provides information about the stability of the system when subject to perturbations of the linear modes, and by adjusting the parameters, we can reduce the conservativeness of the approach to zero. The update of the parameters can be made in an systematic way, thereby providing a semi-complete algorithm for checking the existence of polyhedral Lyapunov functions for piecewise linear systems.

The finite set of points for which the Lyapunov constraints are enforced is obtained in a counterexample-guided fashion. More precisely, at each step of the process, a candidate polyhedral Lyapunov function is computed based on the constraints given by the current set of points. Then, the algorithm checks whether the candidate Lyapunov function is a valid Lyapunov function for the system, and if not, outputs a point, called a counterexample, at which the Lyapunov conditions are violated. This point is added to the set of points at which the constraints are enforced, thereby preventing the candidate to be re-visited by our algorithm.

A desirable property of our approach is that the steps described above are implemented by solving a series of convex optimization problems whose sizes are bounded by the dimension of the system and the number of counterexamples so far. At the same time, we prove that the number of iterations of the process is bounded and derive upper bounds on it (Section 7). We evaluate our approach on a series of numerical examples ranging from challenging instances that have been considered in other works, and a family of piecewise linear systems known to be stable and with dimension from 2 to 9 and number of modes up to 8 (Section 8). We show that our approach terminates faster than the conservative upper bounds established by our theoretical analysis.

2 Comparison with other works

2.1 Polyhedral Lyapunov functions

Compared to Polański (2000); Johansson (2003); Lazar and Doban (2011), our approach does not fix a priori the domain of the linear pieces or the directions of the vertices of the polyhedral Lyapunov function. One can

iteratively solve the problem with increasingly expressive templates, but this might result in overly complex templates if the update of the template is not based on previous computations. Furthermore, unlike Ambrosino et al. (2012); Kousoulidis and Forni (2021); Berger and Sankaranarayanan (2022), our approach does not set a-priori bounds on the number of linear pieces or vertices of the Lyapunov function. Finding a polyhedral Lyapunov function with fixed number of linear pieces or vertices for linear systems amounts to solve a nonconvex optimization problem. Ambrosino et al. (2012); Kousoulidis and Forni (2021) propose convex tightenings or approximate methods to make the computation tractable. One drawback is that, in case of failure, this provides little insight on the stability of the system, whereas our approach allows us to conclude that the system is not stable under small perturbations of the system with predefined bounds. Note that most of the above approaches are restricted to switched linear systems.

Set-theoretic methods (e.g., Miani and Savorgnan, 2005; Guglielmi and Protasov, 2013; Blanchini and Miani, 2015) aim to find an invariant set for (piecewise) linear systems in discrete time by recursively computing the image of an initial polyhedral set by the system until a fixed point is reached. Like our approach, these methods do not place a-priori bounds on the number of linear pieces of the function. However, a clear complexity analysis remains elusive, mainly because of the difficulty of bounding the complexity of the image of polyhedral sets by piecewise linear systems. Guglielmi et al. (2017) extends Guglielmi and Protasov (2013) to continuous-time switched linear systems by using a time discretization of the system to provide lower bounds on the convergence rate of the trajectories. We use a similar approach for piecewise linear systems and show that the time discretization provides a lower bound on the convergence rate of the trajectories under bounded perturbations of the linear modes of the system (Section 5).

2.2 Piecewise polynomial Lyapunov functions

Quadratic Lyapunov functions provide a universal template to study the stability of linear systems (Antsaklis and Michel, 2006). However, they are conservative for switched or piecewise linear systems (Jungers, 2009). This limitation can be alleviated by considering polynomial functions of higher degree (Papachristodoulou and Prajna, 2002) or piecewise quadratic functions (Hassibi and Boyd, 1998; Johansson and Rantzer, 1998; Legat et al., 2020). However, the computational complexity of the polynomial Lyapunov functions grows rapidly with increasing dimension of the system and degree of the polynomial. On the other hand, for piecewise quadratic functions we need to define the domain of the quadratic pieces which may result in a lot of hyper-parameters. Therefore, these techniques are generally restricted to systems of low dimension, although these systems are

not necessarily piecewise linear.

2.3 Data-based Lyapunov analysis

The idea of learning Lyapunov functions from data and verifying the result has received a lot of attention from the control community in recent years. As an early example, Topcu et al. (2008) proposes to learn a candidate Lyapunov function from sampled trajectories and verify the result. Our approach falls more specifically into the category of Counterexample-Guided Inductive Synthesis (CeGIS), which consists in iteratively adding counterexamples from a verification (aka. falsification) step. CeGIS has been used in a wide range of contexts (e.g., Kapinski et al., 2014; Prabhakar and Soto, 2016; Chang et al., 2019; Ravanbakhsh and Sankaranarayanan, 2019; Poonawala, 2021; Ahmed et al., 2020; Abate et al., 2021; Dai et al., 2021). Particularly relevant to our work, Dai et al. (2021) searches for neural-network Lyapunov functions with ReLU activation functions, using Mixed-Integer Linear Programming for the falsification; and Polański (2000); Poonawala (2021) search for polyhedral Lyapunov functions, using Linear Programming for the verification and updating the domain of the linear pieces from the counterexamples. However, both approaches lack guarantees of convergence or complexity. Our work focuses on piecewise linear systems and polyhedral Lyapunov functions. This allows us to avoid the use of SOS relaxations in favor of Linear Programming that lend themselves to precise and efficient solvers. Furthermore, we introduce a “gap” (find a polyhedral Lyapunov function vs conclude that the system is not robustly stable) in our formulation in order to provide formal guarantees of termination. This approach was first introduced in computer science, under the name of δ -*completeness*, to provide practical solutions to problems that are known to be undecidable or intractable (Gao et al., 2012).

In a recent work (Berger and Sankaranarayanan, 2022), we provided a counterexample-guided method to compute polyhedral Lyapunov functions with fixed number of linear pieces. In the present work, we do not place a-priori bounds on the number of linear pieces. This allows us to associate a linear piece to each of the counterexamples, and thereby formulate the problem of candidate learning as a convex optimization problem. This differs from the approach in the previous work, in which we have to solve a Mixed-Integer Program (addressed with a counterexample-guided branch-and-bound approach). Furthermore, the method presented here provides guarantees that if it fails then no polyhedral Lyapunov function satisfying user-specified limits on robustness exists for the underlying system.

3 Notation

$\|\cdot\|$ denotes a vector norm in \mathbb{R}^d (e.g., the L^1 norm), and $\mathbb{S} = \{x \in \mathbb{R}^d : \|x\| = 1\}$ is the associated unit sphere.

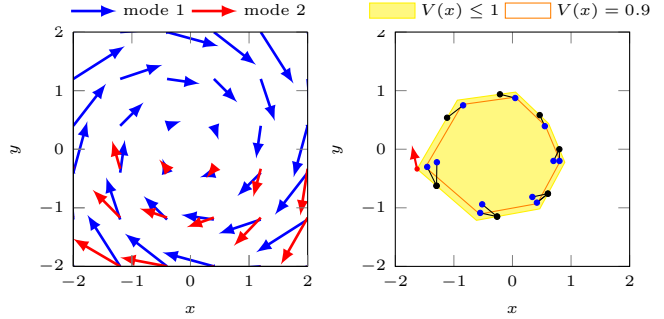


Fig. 1. Piecewise linear dynamics and candidate polyhedral Lyapunov function.

By extension, $\|\cdot\|$ also denotes the matrix norm induced by $\|\cdot\|$ in $\mathbb{R}^{d \times d}$, defined by $\|A\| = \max \{\|Ax\| : x \in \mathbb{S}\}$. $\|\cdot\|_*$ denotes the dual norm of $\|\cdot\|$, defined by $\|c\|_* = \max \{c^T x : x \in \mathbb{S}\}$ (e.g., if $\|\cdot\|$ is the L^1 norm, then $\|\cdot\|_*$ is the L^∞ norm). $\mathbb{B}^* = \{c \in \mathbb{R}^d : \|c\|_* \leq 1\}$ denotes the unit ball associated to $\|\cdot\|$.

All proofs can be found in Appendix A.

4 Problem statement

4.1 Piecewise linear systems

We study piecewise linear dynamical systems in continuous time:

Definition 1. A (continuous-time) *piecewise linear dynamical systems* is a system described by a finite set of modes Q , wherein each mode $q \in Q$ is associated with a region $H_q \subseteq \mathbb{R}^d$ that is a closed polyhedral cone¹, and a transition matrix $A_q \in \mathbb{R}^{d \times d}$. The dynamics is given by the differential inclusion

$$\xi'(t) \in \mathcal{F}(\xi(t)),$$

wherein $\mathcal{F}(x) = \{A_q x : q \in Q, x \in H_q\}$.

The regions H_q are assumed to form a cover of the state space \mathbb{R}^d . However, they may overlap arbitrarily (i.e., they are *not required* to have an intersection of measure 0); if they do not overlap, then the system is deterministic, otherwise, it is uncertain (i.e., trajectories starting from a given point might not be unique).

Note that the set-valued function \mathcal{F} in Definition 1 completely describes the dynamics of the system. Therefore, in the following, we will refer to this system as System \mathcal{F} ; the parameters Q , and H_q and A_q for each $q \in Q$, being implicit in the definition of \mathcal{F} .

¹ That is, $H_q = \{x \in \mathbb{R}^d : M_q x \geq 0\}$ for some $M_q \in \mathbb{R}^{m_q \times d}$ and $m_q \in \mathbb{N}$.

Definition 2. An absolutely continuous function $\xi : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^d$ is a *trajectory* of System \mathcal{F} if $\xi'(t) \in \mathcal{F}(\xi(t))$ for almost all $t \in \mathbb{R}_{\geq 0}$. The system is *asymptotically stable* if all trajectories converge toward the origin.

Example 3 (Running illustrative example). Consider the piecewise linear system described by $Q = \{1, 2\}$, $H_1 = \mathbb{R}^2$, $H_2 = \mathbb{R} \times \mathbb{R}_{\leq 0}$,

$$A_1 = \begin{bmatrix} -0.2 & 1.0 \\ -1.0 & -0.2 \end{bmatrix} \quad \text{and} \quad A_2 = \begin{bmatrix} 0.01 & 1.0 \\ -1.0 & 0.01 \end{bmatrix}.$$

The vector field of the system is represented in Figure 1-left. Note that this system is uncertain. Throughout the paper, we will prove that this system is asymptotically stable by computing a *polyhedral Lyapunov function* for it.

4.2 Polyhedral Lyapunov functions

We aim to study the stability of System \mathcal{F} using Lyapunov analysis. Let us recall that a continuous function $V : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ is a *Lyapunov function* for System \mathcal{F} if (i) $V(x) = 0 \Leftrightarrow x = 0$, (ii) V is radially unbounded, and (iii) for every trajectory ξ with $\xi(0) \neq 0$, V decreases along ξ , i.e., $\forall t > 0, V(\xi(t)) < V(\xi(0))$. It is well known that if System \mathcal{F} admits a Lyapunov function, then it is asymptotically stable (Khalil, 2002, Theorem 4.2).

In this paper, we look for *positively homogeneous convex* functions, also called *gauges*, as Lyapunov functions.

Definition 4. A *gauge* is defined as the pointwise maximum of a (possibly infinite) compact set of linear functions, i.e., $V(x) = \max_{c \in \mathcal{V}} c^\top x$, wherein $\mathcal{V} \subseteq \mathbb{R}^d$ is a compact set of *coefficient vectors*.

With a small abuse of notation, given a gauge V , we let \mathcal{V} be the compact set of its coefficient vectors (note that \mathcal{V} is not uniquely defined, but this will not be an issue in this paper). We also let $\mathcal{V}_{\max} = \max_{c \in \mathcal{V}} \|c\|_*$, and for all $x \in \mathbb{R}^d$, we let $\mathcal{V}(x) = \{c \in \mathcal{V} : V(x) = c^\top x\}$ be the set of coefficient vectors that are maximal at x .

Definition 5. A gauge V for which \mathcal{V} is finite is called a *polyhedral* function.

A *sufficient* condition for a gauge to be a Lyapunov function for System \mathcal{F} is as follows:

Proposition 6. A gauge V is a Lyapunov function for System \mathcal{F} if the following conditions hold:

- (C1) $\forall x \neq 0, V(x) > 0$.
- (C2) $\forall x \neq 0, \forall v \in \mathcal{F}(x), \forall c \in \mathcal{V}(x), c^\top v < 0$.

Although the conditions in Proposition 6 are not necessary for a gauge V to be a Lyapunov function for System \mathcal{F} , they become necessary if we further require that V remains a Lyapunov function for small perturbations of System \mathcal{F} , where perturbations of the regions H_q and the matrices A_q are considered. This will be formalized in the next section.

5 Properties of Lyapunov gauges

5.1 Eccentricity and robustness

We recast the conditions in Proposition 6 in a form that makes explicit two features of Lyapunov gauges (LGs): the *eccentricity* and the *robustness* to system perturbations.

Proposition 7. A gauge V with $\mathcal{V}_{\max} > 0$ satisfies the conditions of Proposition 6 if and only if there are constants $\epsilon \geq 1$ (called the *eccentricity*), $\tau > 0$ (called the *time step*) and $\gamma \in (0, 1)$ (called the *contraction rate*) such that the following conditions hold:

- (D1) $\forall x \in \mathbb{R}^d, V(x) \geq \frac{1}{\epsilon} \mathcal{V}_{\max} \|x\|$.
- (D2) $\forall x \in \mathbb{R}^d, \forall v \in \mathcal{F}(x), V(x + \tau v) \leq \gamma V(x)$.

Remark 8. The smallest value of ϵ satisfying (D1) is equal to the ratio of the largest radius of the 1-level set of V ($\max \{\|x\| : V(x) = 1\}$) by its smallest radius ($\min \{\|x\| : V(x) = 1\}$), hence the name *eccentricity*.²

Combined together, the parameters (ϵ, τ, γ) give a measure of the *robustness* of V as a Lyapunov function with respect to perturbations of the vector field \mathcal{F} , as defined below.

Definition 9 (Perturbed system). Given $\delta_1, \delta_2 \geq 0$, a (δ_1, δ_2) -*perturbation* of System \mathcal{F} is a piecewise linear system \mathcal{F}' with set of modes $Q' = Q$, regions H'_q satisfying $H'_q \subseteq \{x' : \|x' - x\| \leq \delta_1 \|x'\|, x \in H_q\}$ for each $q \in Q$, and matrices A'_q satisfying $\|A'_q - A_q\| \leq \delta_2$ for each $q \in Q$.

Let $\sigma = \max \{\|A_q\| : q \in Q\}$.

Theorem 10 (Sufficient condition for robust LG). *Let V be a gauge satisfying the conditions in Proposition 7. Let $\delta_1, \delta_2 \geq 0$ be such that $(2 + \tau\sigma)\delta_1 + \tau\delta_2 < \frac{1-\gamma}{\epsilon}(1 - \delta_1)$. Then, V is a Lyapunov function for any (δ_1, δ_2) -perturbation of System \mathcal{F} .*

² Notions similar to the eccentricity, but without a proper name, have been used in the literature on quadratic Lyapunov functions to refer to various ratios of the eigenvalues of positive semidefinite matrices (e.g., Kenanian et al., 2019; Berger et al., 2022).

Theorem 11 (Necessary condition for robust LG). *Let $\delta_1, \delta_2 > 0$. Let V be a gauge. Assume that V is a Lyapunov function for any (δ_1, δ_2) -perturbation of System \mathcal{F} . Then, (D2) in Proposition 7 holds with all $\tau \in (0, \frac{1}{\sigma})$ and $\gamma \in (0, 1)$ that satisfy $\frac{\tau\sigma}{1-\tau\sigma} < \delta_1$ and $-\log(1 - \tau\sigma) - \tau\sigma - \log(\gamma) \leq \tau\delta_2$.*

Remark 12. Note that the relations between $\epsilon, \tau, \gamma, \delta_1$ and δ_2 are independent of the dimension of the system and the number of modes. Also, we observe that the parameters ϵ, τ and γ are invariant with respect to positive scaling of V . That is, if V satisfies Proposition 7 with ϵ, τ and γ , then so does the function $\frac{1}{\lambda}V$ for any $\lambda > 0$. Therefore, in the following, we restrict our attention to gauges with $\mathcal{V}_{\max} \leq 1$, i.e., with $\mathcal{V} \subseteq \mathbb{B}^*$.

The relation between the parameters (τ, γ) and the robustness of the Lyapunov function with respect to perturbations of the system being established, we focus in the following of the paper on finding Lyapunov gauges for System \mathcal{F} with robustness parameters (τ, γ) .

5.2 Detection of non-robustness

Given a finite set of points $X \subseteq \mathbb{R}^d$ and parameters ϵ, τ and γ , one can verify whether the conditions in Proposition 7 are satisfied at the points in X by some gauge V . If this is not the case, then one concludes that the system does not admit a Lyapunov gauge with eccentricity ϵ and robustness parameters (τ, γ) .

Proposition 13. *Let $X \subseteq \mathbb{R}^d$ be a finite set of points. Consider the problem of finding a set of coefficient vectors $\mathcal{V}_X = \{c_x : x \in X\} \subseteq \mathbb{B}^*$ (one coefficient vector for each point in X) such that $\forall x \in X, \forall v \in \mathcal{F}(x), \forall c \in \mathcal{V}_X, c_x^\top x \geq \frac{1}{\epsilon}\|x\|, c_x^\top(x + \tau v) \leq \gamma c_x^\top x$ and $\|x + \tau v\| \leq \gamma \epsilon c_x^\top x$. If there exists no such set of coefficient vectors, then one concludes that the system does not satisfy the conditions in Proposition 7 with ϵ, τ and γ for any gauge.*

Remark 14. The problem in Proposition 13 can be formulated as a convex optimization problem with decision variables $\{c_x : x \in X\} \subseteq \mathbb{B}^*$ (in fact, all constraints are linear, except possibly the constraints $c_x \in \mathbb{B}^*$, depending on $\|\cdot\|$). This optimization problem can be solved efficiently and accurately using for instance interior-point algorithms (Boyd and Vandenberghe, 2004).

If the problem in Proposition 13 has a feasible solution \mathcal{V}_X , then the associated gauge V_X will be called a *candidate* Lyapunov gauge for System \mathcal{F} , in the sense that it satisfies the sufficient conditions of Proposition 6 at all points in X , but needs to be *verified* for other points in \mathbb{R}^d . Note that V_X is a *polyhedral function* since \mathcal{V}_X is finite. This procedure of detecting non-robustness and computing a candidate polyhedral Lyapunov function, followed by verifying the candidate, will form the basis of a *counterexample-guided iterative process* to compute

polyhedral Lyapunov functions for piecewise linear systems, described in the next section.

Example 3 (Continued). Consider the system of Example 3, whose vector field is represented in Figure 1-left, and consider the set X consisting in the black dots in Figure 1-right. For $\tau = 0.25$ and each $x \in X$ and $v \in \mathcal{F}(x)$, the points $x + \tau v$ are represented by blue dots in Figure 1-right. The yellow region represents the 1-sublevel set of a candidate polyhedral Lyapunov function satisfying the conditions of Proposition 13 with $\gamma = 0.9$: indeed, we see that the blue dots are inside the γ -sublevel set of V , represented by the orange curve.

6 Algorithm to compute polyhedral Lyapunov functions

The algorithm takes as input three parameters: an eccentricity $\epsilon > 0$, a time step $\tau > 0$ and a contraction rate $\gamma > 0$. The algorithm returns a polyhedral Lyapunov function for the system, or concludes that no Lyapunov gauge satisfying the conditions in Proposition 7 with ϵ, τ and γ exists for the system.

The algorithm is an iterative process that maintains a finite set of points, called *witnesses*: $X = \{x_1, \dots, x_N\} \subseteq \mathbb{R}^d$. The witness set is initialized to \emptyset . Each step iterates between two algorithms in succession:

- Detecting non-robustness and finding a candidate polyhedral Lyapunov function, as explained in Subsection 5.2.
- Verifying the candidate polyhedral Lyapunov function, i.e., verifying whether the conditions in Proposition 6 are satisfied at all points. If the verification succeeds, we have our desired polyhedral Lyapunov function. Otherwise, we get a point (called a *counterexample*) where the candidate fails to satisfy the Lyapunov conditions.

Thus, at the end of each step, there are three possible outcomes: (a) no Lyapunov gauge satisfying the conditions in Proposition 7 with ϵ, τ and γ exists for the system, and the algorithm stops (or update the parameters); (b) the candidate polyhedral function verifies the Lyapunov conditions, and the algorithm stops; or (c) a new witness point is added to X . This process eventually terminates, and we provide upper bounds on the total number of iterations to termination.

In Subsection 5.2, we explained how to efficiently detect non-robustness or compute a candidate polyhedral Lyapunov function. In the next subsection, we explain how to verify this candidate and compute a counterexample if the verification fails.

6.1 Verification and falsification

This step relies on the following proposition, whose proof follows directly from Proposition 6.

Proposition 15. *Let V be a polyhedral Lyapunov function such that $\forall x \neq 0, V(x) > 0$. Consider the problem of finding a mode $q \in Q$, a point $x \in H_q \setminus \{0\}$ and a coefficient vector $c \in \mathcal{V}(x)$ such that $c^\top A_q x \geq 0$. If there exist no such mode, point and coefficient vector, then one concludes that V is a Lyapunov function for the system.*

The above can be addressed as follows: given $q \in Q$ and $c \in \mathcal{V}$, consider the following optimization problem:

$$\begin{aligned} \max \quad & c^\top A_q x \\ \text{s.t.} \quad & x \in H_q \wedge c^\top x = 1 \wedge (\forall c' \in \mathcal{V}) c'^\top x \leq 1. \end{aligned} \quad (1)$$

The problem in Proposition 15 has a solution if and only if there is c and q for which (1) has an optimal solution $x_{c,q}$ satisfying $c^\top A_q x_{c,q} \geq 0$. In this case, we choose as *counterexample* the point $x_{c,q}$ for which $c^\top A_q x_{c,q}$ is the largest among all c and q .

Remark 16. Solving (1) amounts to solve a linear programs with d variables and $|\mathcal{V}| + m_q$ constraints (where m_q is the number of linear constraints describing H_q). This can be done very efficiently and reliably using Linear Programming solvers.

It remains to explain how we ensure that the polyhedral function V that is verified satisfies $\forall x \neq 0, V(x) > 0$. To do this, we fix $\eta \in (0, \epsilon)$, and let $\mathcal{V}_o \subseteq \frac{1}{\epsilon} \mathbb{B}^*$ be a finite set of coefficient vectors such that $\forall x \in \mathbb{R}^d, V_o(x) \geq \eta \|x\|$. Then, given a candidate polyhedral Lyapunov function V_X , we let V be defined by $\mathcal{V} = \mathcal{V}_X \cup \mathcal{V}_o$. This ensures that $\forall x \in \mathbb{R}^d, V(x) \geq V_o(x) \geq \eta \|x\|$.

Remark 17. Depending on the norm $\|\cdot\|$ and the ratio η/ϵ , the cardinality of \mathcal{V}_o may vary;³ but in any case, this set is fixed through all iterations of the algorithm.

Example 3 (Continued). Consider the system of Example 3, whose vector field is represented in Figure 1-left, and the candidate polyhedral Lyapunov function whose 1-sublevel set is represented in Figure 1-right. The red dot in Figure 1-right is a counterexample found by solving (1): indeed, we see that the flow direction according to mode 2 is toward the exterior of the sublevel set.

6.2 Overall algorithm

The overall algorithm is described in Algorithm 1. The process starts with an empty set of witnesses. Then, it

³ When $\|\cdot\|$ is the L^∞ -norm, we can choose $\eta = \epsilon$ and \mathcal{V}_o with cardinality $2d$.

Algorithm 1: Learning a polyhedral Lyapunov function.

Data: System \mathcal{F} , eccentricity $\epsilon > 0$, time step $\tau > 0$ and contraction rate γ .

Result: Polyhedral Lyapunov function V , or FAIL.

Let $X_0 \leftarrow \emptyset$

for $k = 0, 1, \dots$ **do**

Compute a candidate polyhedral Lyapunov function V_k by solving the problem in Proposition 13 with $X = X_k$

if infeasible then return FAIL

Find a counterexample x_k for the candidate V_k by solving the problem in Proposition 15 with $\mathcal{V} = \mathcal{V}_k \cup \mathcal{V}_o$.

if infeasible then return V_k

Let $X_{k+1} \leftarrow X_k \cup \{x_k\}$

enters a loop, in which, at each iteration, the steps described in Subsections 5.2 and 6.1 are performed sequentially: (i) from the current set of witnesses, it tries to find a candidate polyhedral Lyapunov function for the system with given eccentricity and robustness parameters. If this is not feasible, the algorithm stops and outputs FAIL; (ii) it checks whether the candidate function provides a valid Lyapunov function for the system. If it is the case, then the algorithm stops and outputs the candidate function. Otherwise, it produces a counterexample, which is added to the witness set. The algorithm then proceeds with the next iteration of the loop.

7 Analysis of the algorithm

7.1 Termination and complexity

The algorithm is sound in the sense that if it terminates and outputs a polyhedral function, then this function is a Lyapunov function for the system (Proposition 15); otherwise, if it outputs FAIL, then the system does not admit a Lyapunov gauge with eccentricity ϵ and robustness parameters (τ, γ) (Proposition 13).

We show that the algorithm terminates and we bound the number of steps to termination. The proof of termination exploits the gap between the constraints that are enforced during the candidate generation phase and the constraints that are checked during the verification phase: namely, at the witness points, the constraints enforced during the generation phase are stricter (robust Lyapunov constraints) than the constraints checked at these points during the verification phase. This implies that the produced counterexample cannot be part of the current witness set, and furthermore, its distance to the witness set, after normalization, is bounded from below by a positive constant that depends on the system and the parameters ϵ, τ and γ .

Definition 18. Let $X \subseteq X' \subseteq \mathbb{R}^d \setminus \{0\}$ and $r \geq 0$. We

say that X' is an r -inflation of X if $\forall q \in Q$ and $\forall x' \in X' \cap H_q$, there is $x \in X \cap H_q$ such that $\left\| \frac{x}{\|x\|} - \frac{x'}{\|x'\|} \right\| < r$.

Let $r = \frac{1-\gamma}{(2+\tau\sigma)\epsilon}$, where $\sigma = \max\{\|A_q\| : q \in Q\}$.

Proposition 19. *Let $X \subseteq \mathbb{R}^d \setminus \{0\}$. Let V_X be a solution to the problem in Proposition 13 and $\mathcal{V} = \mathcal{V}_X \cup \mathcal{V}_\circ$. Let $X' \subseteq \mathbb{R}^d \setminus \{0\}$ be an r -inflation of X . Then, $\forall x' \in X'$, $\forall v' \in \mathcal{F}(x')$ and $\forall c' \in \mathcal{V}(x')$, $c'^T v' < 0$.*

Corollary 20. *Let $X_k \subseteq X_{k+1}$ be two consecutive witness sets generated during the execution of Algorithm 1. Then, X_{k+1} is not an r -inflation of X_k .*

The proof is direct from Proposition 19 and the definition of $X_{k+1} = X_k \cup \{x_k\}$, wherein the counterexample x_k violates the condition of Proposition 15.

From Corollary 20, we derive the following upper bound on the number of iterations of the algorithm. For every $s > 0$, let $\text{Pack}(s; \mathbb{S})$ denote the s -packing number of \mathbb{S} , i.e. the largest cardinality of a subset $S \subseteq \mathbb{S}$ such that $\forall x \in S, \forall y \in S, x \neq y$ implies $\|x - y\| \geq s$.

Theorem 21 (Termination). *Algorithm 1 terminates in at most $|Q| \text{Pack}(r; \mathbb{S})$ steps.*

The overall arithmetic complexity of Algorithm 1 is thus in $\mathcal{O}(d^\alpha (|Q| \text{Pack}(r; \mathbb{S}))^{\alpha+2\beta+1})$, when using a convex optimization solver whose complexity is in $\mathcal{O}(n^\alpha m^\beta)$ with n the number of variables and m the number of constraints.⁴ Note that $\text{Pack}(r; \mathbb{S})$ grows as $\mathcal{O}(r^{1-d})$.⁵ This is an upper bound on the worst-case complexity; as we will see in the next section, the algorithm performs much more efficiently than this bound on several practical examples.

7.2 Choice of the parameters

The choice of the eccentricity ϵ can be driven by safety considerations: e.g., one may want to ensure that if the system starts in a state of norm 1, then the trajectories do not diverge to a state of norm larger than ϵ before converging toward the origin. The choice of the robustness parameters (τ, γ) are driven by robustness considerations with respect to system perturbations: e.g., if one believe that the system is stable under (δ_1, δ_2) -perturbations, this provides lower bounds on the value of (τ, γ) according to Proposition 11.

If the value of (δ_1, δ_2) is unknown, one can use the following strategy to update (τ, γ) starting from some initial guess (τ_0, γ_0) : $(\tau_{\ell+1}, \gamma_{\ell+1}) = (\tau_\ell / \sqrt{\alpha}, 1 - (1 - \gamma_\ell) / \alpha)$,

⁴ For Linear Programs, $\alpha = 2$ and $\beta = 1.5$ using interior-point methods (Ben-Tal and Nemirovski, 2001, p. 422).

⁵ For the L^∞ -norm, $\text{Pack}(r; \mathbb{S}) \leq (2r^{-1} + 1)^{d-1} 2d$.

where $\alpha > 1$. One can check that, with this strategy, (τ_ℓ, γ_ℓ) will eventually satisfy the conditions in Proposition 11, provided $\delta_1, \delta_2 > 0$. Per Proposition 21, the number of iterations of the algorithm increases by a factor α^{d-1} at each update of (τ_ℓ, γ_ℓ) . If ϵ is not fixed a priori neither, one can update it as $\epsilon_{\ell+1} = \alpha \epsilon_\ell$, thereby increasing the number of iterations of the algorithm by a factor $\alpha^{2(d-1)}$ at each update of $(\epsilon_\ell, \tau_\ell, \gamma_\ell)$. Finally, note that, after an update of the parameters, we do not need to reset the witness set; however, it is not clear whether reusing the previous witness set results in a significant speed-up of the algorithm.

8 Numerical experiments

We use the L^∞ -norm for $\|\cdot\|$. The problems in Propositions 13 and 15 can then be formulated as Linear Programs. All computations were made on a laptop with processor Intel Core i7-7600u and 16GB RAM running Windows. We used GurobiTM 10.0, under academic license, as linear optimization solver.

8.1 Example 3 (finished)

Consider the system of Example 3. We want to show that the system is asymptotically stable and that the trajectories starting with norm $\rho \geq 0$ never reach a state with norm larger than 3ρ . Therefore, we fix $\epsilon = 3$. Then, to find parameters (τ, γ) that allows us to prove stability of the system, we use the strategy described in Subsection 7.2 with $(\tau_0, \gamma_0) = (1, 1)$ and $\alpha = 4$. After two updates of the parameters, i.e., with $\tau_2 = 1/4$ and $\gamma_2 = 1$, Algorithm 1 finds a polyhedral Lyapunov function for the system in 32 steps (the steps are illustrated in Figure 2).

In the process of updating (τ, γ) , we also learn that the system does not admit a polyhedral Lyapunov function with parameters $(\epsilon, \tau, \gamma) = (3, 1/2, 1)$ (otherwise, the algorithm would have found a polyhedral Lyapunov function after one update of the parameters). Per Proposition 11, this implies that there exists a (δ_1, δ_2) -perturbation of the system, with $\delta_1 < 1.5$ and $\delta_2 < 0.633$ for which there is a trajectory starting with norm 1 that reaches a state with norm larger than 3.

8.2 Benchmark: 2D uncertain linear system

This system, introduced by Zelentsovsky (1994), is described by $\mathcal{F}(x) = \{A_p x : p \in \{0, \alpha\}\}$ where

$$A_p : \begin{bmatrix} 0 & 1 \\ -2 & -1 \end{bmatrix} + p \begin{bmatrix} 0 & 0 \\ -1 & 0 \end{bmatrix}.$$

Zelentsovsky (1994) shows that the system with $\alpha = 3.82$ admits a quadratic Lyapunov function. Blanchini

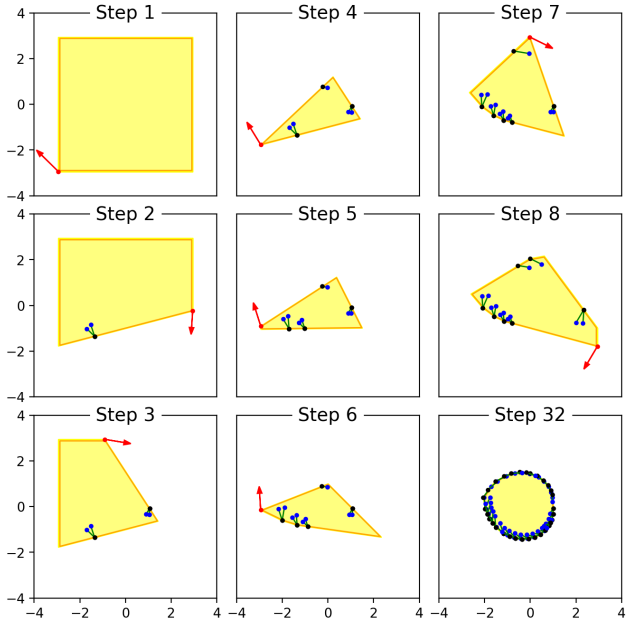


Fig. 2. Steps of the construction of a polyhedral Lyapunov function for the system of Example 3, with $\epsilon = 3$, $\tau = 1/4$ and $\gamma = 1$. At each step $k = 1, \dots, 32$, a polyhedral function V_k (yellow region), satisfying the conditions of Proposition 13 at the witness set X_k (black dots), is computed. Then, the algorithm checks whether it can find a counterexample x_k (red dots). After 32 steps, the process has computed a polyhedral Lyapunov function for the system (last plot).

and Miani (1996) provides a polyhedral Lyapunov function for the system with $\alpha = 6$. Xie et al. (1997) provides a piecewise quadratic function for the system with $\alpha = 6.2$. Chesi et al. (2009) provides a polynomial Lyapunov function of degree 20 for the system with $\alpha = 6.8649$. Ambrosino et al. (2012) provides a polyhedral Lyapunov function with 9694 vertices for the system with $\alpha = 6.87$.

Using Algorithm 1, we computed a polyhedral Lyapunov function for the system with $\alpha = 6$. For that, we fixed $\epsilon = 50$ and updated the parameters (τ, γ) using the strategy described in Subsection 7.2 with $(\tau_0, \gamma_0) = (1, 1)$ and $\alpha = 4$. This led to the feasible parameters $\tau_6 = 1/64$ and $\gamma_6 = 1$. The computation took about 4 minutes (224 iterations), and found a polyhedral Lyapunov functions with 228 linear pieces (see Figure 3a).

Finally, we also applied our algorithm on the system with $\alpha = 6.87$. The computation reached the time-out limit, set to 4 hours, without finding a polyhedral Lyapunov. Nevertheless, we found that the system does not admit a polyhedral Lyapunov function with parameters $(\epsilon, \tau, \gamma) = (50, 1/64, 1)$. From this, we can deduce information about the stability of the system under perturbations of the matrices by using Theorem 11.

Remark 22. Let us mention that the algorithms in the above papers focus on uncertain linear systems, while

our algorithm also tackles *piecewise* linear systems. Another difference, e.g., with Ambrosino et al. (2012), is that the hyper-parameters in our algorithm are directly related to meaningful properties of the system, namely the existence of a Lyapunov gauge robust to system perturbations.

8.3 Controlled mass–spring system

We consider a mass–spring system whose dynamics is described by $\dot{x} = -20x + u$. We control this system with a PID controller defined by $u(t) = -K_i y(t) - K_p x(t) - K_d \dot{x}(t)$, wherein $y(t) = \int_0^t x(s) ds$, $K_i = 440$, $K_p = 240$ and $K_d = 32$. The force that is applied on the mass can only be *nonnegative*. To counterbalance the accumulation of the error when the input is negative, we add an *anti-windup* mechanism to the system. The block-diagram of the resulting system is depicted in Figure 3c. The dynamics of the system is described by the following piecewise linear system:

$$\dot{y} = x, \quad \ddot{x} + K_d \dot{x} + (20 + K_p)x + K_i y = 0, \quad (2)$$

$$\text{if } K_d \dot{x} + K_p x + K_i y \leq 0;$$

$$\dot{y} = -10y + x, \quad \ddot{x} + 20x = 0, \quad (3)$$

$$\text{if } K_d \dot{x} + K_p x + K_i y \geq 0.$$

Remark 23. Since (3) is not stable, the system does not admit a Lyapunov function symmetric around the origin (including any polynomial Lyapunov function).

Using Algorithm 1, we computed a polyhedral Lyapunov function for this system. We used the parameters $\epsilon = 50$, $\tau = 1/32$ and $\gamma = 1$ (obtained by using the strategy described in Subsection 7.2). The computation took 15 seconds (130 iterations) and found a polyhedral Lyapunov function with 136 linear pieces (see Figure 3b).


8.4 Performance evaluation

We evaluate the performance of the process, in terms of computation time and complexity of the outputted Lyapunov function, as a function of the dimension of the system and the stability margin of the linear modes.


Therefore, for $d \in \{4, 5, \dots, 9\}$, we let $U \in \mathbb{R}^{d \times d}$ be a randomly generated orthogonal matrix and define

$$\Pi = U \mathbf{1} \mathbf{1}^\top U^\top, \quad \mathbf{1} = [1, \dots, 1]^\top \in \mathbb{R}^d.$$

Then, for each $\eta \in \{0.5, 0.05\}$ and $m \in \{1, 2, 4\}$, we define a system with $2m$ modes as follows. For each $i \in \{1, \dots, m\}$, we generate a random vector $a \in \mathbb{S}$ and define the mode $q = 2i - 1$ by $H_q = \{x \in \mathbb{R}^d : a^\top x \geq 0\}$ and $A_q = \Pi - (d + 1)I$, and define the mode $q = 2i$ by $H_q = \{x \in \mathbb{R}^d : a^\top x \leq 0\}$ and $A_q = \Pi - (d + \eta)I$. We fixed $\epsilon = 10$, $\tau = 1/8$ and $\gamma = 1$. We used Algorithm 1 to



(a) Polyhedral Lyapunov function for the system in Subsection 8.2 (b) Polyhedral Lyapunov function for the system in Subsection 8.3



(c) Block-diagram of the mass–spring system in Subsection 8.3 actuated by a ReLU-saturated PID controller. When the actuation force is saturated, an anti-windup mechanism (in blue) counterbalances the accumulation of the error.

Fig. 3.

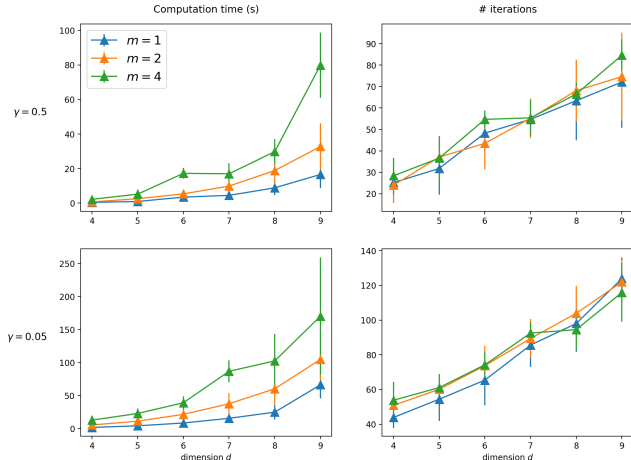


Fig. 4. Average computation time and number of iterations over 10 samples with randomly generated systems, as described in Subsection 8.4. The vertical bars represent the standard deviation over the 10 samples.

compute a polyhedral Lyapunov function for the system with these parameters. For each value of (d, m, η) , we generated 10 different systems and measured the computation time and number of iterations of the algorithm. The results are reported in Figure 4.

9 Extensions and future work

The approach can be readily extended to discrete-time piecewise linear systems. The differences are that a time step parameter τ would not be needed anymore and the verification part (Proposition 15) would account for the conditions of being a Lyapunov function for discrete-time systems. Extension to hybrid linear systems is also possible but the conditions that we impose on the function to be a Lyapunov function might be conservative (e.g., if the discrete transitions include the identity map, do we impose a strict decrease of the Lyapunov function with respect to this map?). One way to tackle this is to add dwell-time assumptions on the system; this is an approach that we will consider in future work, along with extensions to multiple Lyapunov functions.

In future work, we also plan to extend the approach to piecewise affine systems and problems of safety verification, using invariant polytopes. The following two differences with the current approach are to be addressed: first, we can no longer assume that the counterexamples are on the unit sphere; second, the property of decrease is replaced by a property of invariance, in which only the points on the boundary of the polytope must be pushed toward the interior of the polytope.

Acknowledgements

We thank the anonymous reviewers for their insightful comments and suggestions.

References

Alessandro Abate, Daniele Ahmed, Mirco Giacobbe, and Andrea Peruffo. Formal synthesis of Lyapunov neural networks. *IEEE Control Systems Letters*, 5(3):773–778, 2021. doi: 10.1109/LCSYS.2020.3005328.

Amir Ali Ahmadi and Raphaël M Jungers. Lower bounds on complexity of Lyapunov functions for switched linear systems. *Nonlinear Analysis: Hybrid Systems*, 21:118–129, 2016. doi: 10.1016/j.nahs.2016.01.003.

Daniele Ahmed, Andrea Peruffo, and Alessandro Abate. Automated and sound synthesis of Lyapunov functions with SMT solvers. In Armin Biere and David Parker, editors, *International Conference on Tools and Algorithms for the Construction and Analysis of Systems. TACAS 2020.*, volume 12078 of *Lecture Notes in Computer Science*, pages 97–114. Springer, 2020. doi: 10.1007/978-3-030-45190-5_6.

Roberto Ambrosino, Marco Ariola, and Francesco Amato. A convex condition for robust stability analysis via polyhedral Lyapunov functions. *SIAM Journal on Control and Optimization*, 50(1):490–506, 2012. doi: 10.1137/110796479.

- Panos J Antsaklis and Anthony N Michel. *Linear systems*. Birkhäuser, Boston, MA, 2nd edition, 2006. doi: 10.1007/0-8176-4435-0.
- Aharon Ben-Tal and Arkadi Nemirovski. *Lectures on modern convex optimization: analysis, algorithms, and engineering applications*. SIAM, Philadelphia, PA, 2001. doi: 10.1137/1.9780898718829.
- Guillaume O Berger and Sriram Sankaranarayanan. Learning fixed-complexity polyhedral Lyapunov functions from counterexamples, 2022. arXiv preprint arXiv:2204.06693.
- Guillaume O Berger, Raphaël M Jungers, and Zheming Wang. Data-driven invariant subspace identification for black-box switched linear systems, 2022. arXiv preprint arXiv:2209.05320.
- Franco Blanchini and Stefano Miani. On the transient estimate for linear systems with time-varying uncertain parameters. *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, 43(7):592–596, 1996. doi: 10.1109/81.508180.
- Franco Blanchini and Stefano Miani. *Set-theoretic methods in control*. Birkhäuser, Cham, 2nd edition, 2015. doi: 10.1007/978-3-319-17933-9.
- Vincent D Blondel and John N Tsitsiklis. Complexity of stability and controllability of elementary hybrid systems. *Automatica*, 35(3):479–489, 1999. doi: 10.1016/S0005-1098(98)00175-7.
- Stephen Boyd and Lieven Vandenbergh. *Convex optimization*. Cambridge University Press, Cambridge, UK, 2004. doi: 10.1017/CBO9780511804441.
- Ya-Chien Chang, Nima Roohi, and Sicun Gao. Neural Lyapunov control. In *NIPS’19: Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pages 3245–3254. ACM, 2019. doi: 10.5555/3454287.3454579.
- Graziano Chesi, Andrea Garulli, Alberto Tesi, and Antonio Vicino. *Homogeneous polynomial forms for robustness analysis of uncertain systems*, volume 390 of *Lecture Notes in Control and Information Sciences*. Springer, London, 2009. doi: 10.1007/978-1-84882-781-3.
- Frank J Christophersen. Piecewise affine systems. In *Optimal control of constrained piecewise affine systems*, volume 359 of *Lecture Notes in Control and Information Sciences*, pages 39–42. Springer, Berlin, 2007. doi: 10.1007/978-3-540-72701-9_4.
- Hongkai Dai, Benoit Landry, Lujie Yang, Marco Pavone, and Russ Tedrake. Lyapunov-stable neural-network control. In *Proceedings of Robotics: Science and Systems*, Virtual, 2021. doi: 10.15607/RSS.2021.XVII.063.
- Sicun Gao, Jeremy Avigad, and Edmund M Clarke. δ -complete decision procedures for satisfiability over the reals. In Bernhard Gramlich, Dale Miller, and Uli Sattler, editors, *6th International Joint Conference on Automated Reasoning, IJCAR 2012*, volume 7364 of *Lecture Notes in Computer Science*, pages 286–300, Berlin, 2012. Springer. doi: 10.1007/978-3-642-31365-3_23.
- Gene H Golub and Charles F Van Loan. *Matrix computations*. The Johns Hopkins University Press, Baltimore, MD, 4th edition, 2013.
- Nicola Guglielmi and Vladimir Protasov. Exact computation of joint spectral characteristics of linear operators. *Foundations of Computational Mathematics*, 13(1):37–97, 2013. doi: 10.1007/s10208-012-9121-0.
- Nicola Guglielmi, Linda Laglia, and Vladimir Protasov. Polytope Lyapunov functions for stable and for stabilizable LSS. *Foundations of Computational Mathematics*, 17(2):567–623, 2017. doi: 10.1007/s10208-015-9301-9.
- Arash Hassibi and Stephen Boyd. Quadratic stabilization and control of piecewise-linear systems. In *Proceedings of the 1998 American Control Conference. ACC (IEEE Cat. No. 98CH36207)*, pages 3659–3664. IEEE, 1998. doi: 10.1109/ACC.1998.703296.
- Mikael Johansson. *Piecewise linear control systems: a computational approach*, volume 284 of *Lecture Notes in Control and Information Sciences*. Springer, Berlin, 2003. doi: 10.1007/3-540-36801-9.
- Mikael Johansson and Anders Rantzer. Computation of piecewise quadratic Lyapunov functions for hybrid systems. *IEEE Transactions on Automatic Control*, 43(4):555–559, 1998. doi: 10.1109/9.664157.
- Raphaël M Jungers. *The joint spectral radius: theory and applications*. Springer, Berlin, 2009. doi: 10.1007/978-3-540-95980-9.
- James Kapinski, Jyotirmoy V Deshmukh, Sriram Sankaranarayanan, and Niko Aréchiga. Simulation-guided Lyapunov analysis for hybrid dynamical systems. In *Proceedings of the 17th international conference on Hybrid systems: computation and control*, pages 133–142. ACM, 2014. doi: 10.1145/2562059.2562139.
- Joris Kenanian, Ayca Balkan, Raphaël M Jungers, and Paulo Tabuada. Data driven stability analysis of black-box switched linear systems. *Automatica*, 109:108533, 2019. doi: 10.1016/j.automatica.2019.108533.
- Hassan K Khalil. *Nonlinear systems*. Prentice-Hall, Upper Saddle River, NJ, 3rd edition, 2002.
- Dimitris Kousoulidis and Fulvio Forni. Polyhedral Lyapunov functions with fixed complexity. In *2021 60th IEEE Conference on Decision and Control (CDC)*, pages 3293–3298. IEEE, 2021. doi: 10.1109/CDC45484.2021.9683697.
- Mircea Lazar and Alina I Doban. On infinity norms as Lyapunov functions for continuous-time dynamical systems. In *2011 50th IEEE Conference on Decision and Control and European Control Conference*, pages 7567–7572. IEEE, 2011. doi: 10.1109/CDC.2011.6161163.
- Benoît Legat, Saša V Raković, and Raphaël M Jungers. Piecewise semi-ellipsoidal control invariant sets. *IEEE Control Systems Letters*, 5(3):755–760, 2020. doi: 10.1109/LCSYS.2020.3005326.
- Stefano Miani and Carlo Savorgnan. MAXIS-G: a software package for computing polyhedral invariant sets for constrained LPV systems. In *Proceedings of*

the 44th IEEE Conference on Decision and Control, pages 7609–7614. IEEE, 2005. doi: 10.1109/CDC.2005.1583390.

Antonios Papachristodoulou and Stephen Prajna. On the construction of Lyapunov functions using the sum of squares decomposition. In *Proceedings of the 41st IEEE Conference on Decision and Control, 2002*, pages 3482–3487. IEEE, 2002. doi: 10.1109/CDC.2002.1184414.

Andrzej Polański. On absolute stability analysis by polyhedral Lyapunov functions. *Automatica*, 36(4):573–578, 2000. doi: 10.1016/S0005-1098(99)00180-6.

Hasan A Poonawala. Stability analysis of conewise affine dynamical systems using conewise linear Lyapunov functions. *IEEE Control Systems Letters*, 5(6):2126–2131, 2021. doi: 10.1109/LCSYS.2020.3046702.

Pavithra Prabhakar and Miriam Garcia Soto. Counterexample guided abstraction refinement for stability analysis. In Swarat Chaudhuri and Azadeh Farzan, editors, *Computer Aided Verification. CAV 2016.*, volume 9779 of *Lecture Notes in Computer Science*, pages 495–512, Cham, 2016. Springer. doi: 10.1007/978-3-319-41528-4_27.

Hadi Ravanbakhsh and Sriram Sankaranarayanan. Learning control Lyapunov functions from counterexamples and demonstrations. *Autonomous Robots*, 43(2):275–307, 2019. doi: 10.1007/s10514-018-9791-9.

Zhendong Sun and Shuzhi Sam Ge. *Stability theory of switched dynamical systems*. Springer, London, 2011. doi: 10.1007/978-0-85729-256-8.

Ufuk Topcu, Andrew Packard, and Peter Seiler. Local stability analysis using simulations and sum-of-squares programming. *Automatica*, 44(10):2669–2675, 2008. doi: 10.1016/j.automatica.2008.03.010.

Lin Xie, Serge Shishkin, and Minyue Fu. Piecewise Lyapunov functions for robust stability of linear time-varying systems. *Systems & Control Letters*, 31(3):165–171, 1997. doi: 10.1016/S0167-6911(97)00027-3.

Jun Xu and Lihua Xie. Piecewise affine systems. In *Control and estimation of piecewise affine systems*, pages 17–39. Woodhead, Cambridge, UK, 2014. doi: 10.1533/9781782421627_17.

AL Zelentsovsky. Nonquadratic Lyapunov functions for robust stability analysis of linear uncertain systems. *IEEE Transactions on Automatic Control*, 39(1):135–138, 1994. doi: 10.1109/9.273350.

A Proofs

PROOF of Proposition 6 (C1) implies that $V(x) = 0 \Leftrightarrow x = 0$. Furthermore, since V is positively homogeneous, it follows that V is radially unbounded. Now, let ξ be a trajectory with $\xi(0) \neq 0$. We show the following:

Main result: There is $\alpha > 0$ such that $\forall x \neq 0, \forall v \in \mathcal{F}(x), \forall c \in \mathcal{V}, c^\top x \geq V(x) - \alpha \|x\|$ implies $c^\top v < 0$.

Proof of the main result. For a proof by contradiction, assume that it is not the case. Then, there is a sequence $\{(x_k, c_k, q_k)\}_{k=0}^\infty \subseteq \mathbb{S} \times \mathcal{V} \times Q$ such that $\forall k \in \mathbb{N}, c_k^\top x \geq V(x_k) - 2^{-k}, x_k \in H_{q_k}$ and $c_k^\top A_{q_k} x_k \geq 0$. By compactness of $\mathbb{S} \times \mathcal{V}$ and finiteness of Q , and taking a subsequence if necessary, there is $(x, c, q) \in \mathbb{S} \times \mathcal{V} \times Q$ such that $(x_k, c_k) \rightarrow (x, c)$ and $\forall k \in \mathbb{N}, q_k = q$. Since H_q is closed, it holds that $x = \lim_k x_k \in H_q$. By continuity, it holds that $c^\top x - V(x) = \lim_k c_k^\top x_k - V(x_k) \geq 0$, so $c \in \mathcal{V}(x)$. Finally, it holds that $c^\top A_q x = \lim_k c_k^\top A_q x_k \geq 0$. This is a contradiction with (C2), concluding the proof of the main result.

Let $\alpha > 0$ be as in the main result above. By using the continuity of ξ , let $T > 0$ be such that $\forall t \in [0, T]$, (i) $\xi(t) \neq 0$ (this is in fact true for all $t \geq 0$ but we do not need to prove this), and (ii) $\forall c \in \mathcal{V}(\xi(T)), c^\top \xi(t) \geq V(\xi(t)) - \alpha \|\xi(t)\|$. We show that $V(\xi(T)) < V(\xi(0))$. Therefore, fix $c \in \mathcal{V}(\xi(T))$. Then, by the assumption on T and from the main result, it follows that $\forall t \in [0, T], \forall v \in \mathcal{F}(\xi(t)), c^\top v < 0$. Hence, by definition of ξ being a trajectory, $\int_0^T c^\top \xi'(t) dt < 0$, so that $c^\top \xi(T) < c^\top \xi(0)$. Since $c^\top \xi(T) = V(\xi(T))$ and $c^\top \xi(0) \leq V(\xi(0))$, we get that $V(\xi(T)) < V(\xi(0))$. Note that $T > 0$ can be chosen independently of $\xi(0)$. Hence, the same argument can be applied for a trajectory starting at $\xi(T)$. We conclude that $\forall t > 0, V(\xi(t)) < V(\xi(0))$. \square

PROOF of Proposition 7 For the “if” direction, it is straightforward to see that (D1)–(D2) implies (C1)–(C2) in Proposition 6.

To prove the “only if” direction, assume that V satisfies (C1)–(C2) in Proposition 6. Note that V is continuous since \mathcal{V} is bounded.

First, we show that there is $\epsilon \geq 1$ such that (D1) holds. By (C1) and V being continuous, there is $\epsilon \geq 1$ such that $\forall x \in \mathbb{S}, V(x) \geq \frac{V_{\max}}{\epsilon}$. Since V is positively homogeneous of degree 1, it follows that $\forall x \in \mathbb{R}^d, V(x) \geq \frac{V_{\max}}{\epsilon} \|x\|$.

It remains to show that there are constants $\tau > 0$ and $\gamma \in (0, 1)$ such that (D2) holds. For a proof by contradiction, assume that it is not the case. Then, there is a sequence $\{(x_k, c_k, q_k)\}_{k=0}^\infty \subseteq \mathbb{S} \times \mathcal{V} \times Q$ such that $\forall k \in \mathbb{N}, x_k \in H_{q_k}$ and $c_k^\top (x_k + 2^{-k} A_{q_k} x_k) \geq (1 - 4^{-k}) V(x_k)$. By compactness of $\mathbb{S} \times \mathcal{V}$ and finiteness of Q , and taking a subsequence if necessary, there is $(x, c, q) \in \mathbb{S} \times \mathcal{V} \times Q$ such that $(x_k, c_k) \rightarrow (x, c)$ and $\forall k \in \mathbb{N}, q_k = q$. Since H_q is closed, it holds that $x = \lim_k x_k \in H_q$. By continuity, it holds that $c^\top x - V(x) = \lim_k c_k^\top x_k - V(x_k) \geq \lim_k -4^{-k} V(x_k) - c_k^\top 2^{-k} A_q x_k = 0$, so $c \in \mathcal{V}(x)$. Furthermore, $\forall k \in \mathbb{N}, c_k^\top A_q x_k \geq -2^{-k} V(x_k)$, since $c_k^\top x_k \leq V(x_k)$. Hence, by continuity, $c^\top A_q x = \lim_k c_k^\top A_q x_k \geq \lim_k -2^{-k} V(x_k) = 0$. This is a contradiction with (C2), since $c \in \mathcal{V}(x)$, concluding the proof. \square

PROOF of Theorem 10 Let \mathcal{F}' be a (δ_1, δ_2) -perturbation. Fix $q \in Q$, $x' \in H'_q \cap \mathbb{S}$ and $c \in \mathcal{V}(x')$. Let $x \in H_q$ be such that $\|x' - x\| \leq \delta_1$. Denote $v' = A'_q x'$ and $v = A_q v$. It holds that $\|v' - v\| \leq \|A_q(x' - x)\| + \|(A'_q - A_q)x'\| \leq \sigma\delta_1 + \delta_2$. Also, $|c^\top(x' - x)| \leq \mathcal{V}_{\max}\delta_1$ and $|V(x') - V(x)| \leq \mathcal{V}_{\max}\delta_1$. Finally, $\|x\| \geq 1 - \delta_1$. Thus, $c^\top \tau v' = c^\top(x' + \tau v') - V(x') \leq c^\top(x + \tau v) - V(x) + \mathcal{V}_{\max}(2\delta_1 + \tau\sigma\delta_1 + \tau\delta_2) < c^\top(x + \tau v) - V(x) + \frac{1-\gamma}{\epsilon}(1 - \delta_1)\mathcal{V}_{\max} \leq 0$, where the before-last inequality comes from the assumption on (δ_1, δ_2) and the last inequality from (D1)–(D2). Thus, (C2) in Proposition 6 is satisfied for x' , q and c . Since x' , q and c were arbitrary, this concludes the proof. \square

Lemma 24. Let $B \in \mathbb{R}^{d \times d}$ be such that $\|B - I\| \leq \alpha < 1$. Then, (i) $\log(B) = \sum_{k=1}^{\infty} (-1)^{k+1} \frac{(B-I)^k}{k}$ is well defined and satisfies $e^{\log(B)} = B$; (ii) $\forall t \in [0, 1]$, $\|e^{t \log(B)} - I\| \leq \frac{\alpha}{1-\alpha}$; and (iii) $\|\log(B) - B + I\| \leq -\log(1 - \alpha) - \alpha$.

PROOF. (i) See Golub and Van Loan (2013, p. 541). (ii) From $e^A - I = \sum_{k=1}^{\infty} \frac{A^k}{k!}$, we get that $\|e^A - I\| \leq e^{\|A\|} - 1$. Similarly, we get that $\|\log(B)\| \leq -\log(1 - \alpha)$. Thus, $\|e^{t \log(B)} - I\| \leq e^{t\|\log(B)\|} - 1 \leq e^{\|\log(B)\|} - 1 \leq \frac{1}{1-\alpha} - 1$. (iii) Similar to (ii). \square

PROOF of Theorem 11 Let (τ, γ) be as in the theorem. Fix $q \in Q$ and $x \in H_q \cap \mathbb{S}$. Denote $A = A_q$ and $\alpha = \|\tau A\| \leq \tau\sigma$. Let $A' = \frac{1}{\tau} \log(I + \tau A)$ which exists since $\alpha < 1$ (Lemma 24i and assumption on τ) and satisfies $\|A' - A\| \leq \frac{1}{\tau}(-\log(1 - \alpha) - \alpha)$ (Lemma 24iii). Let $R(x) = \{e^{tA'}x : t \in [0, \tau]\}$. By Lemma 24ii, it holds that $\forall y \in R(x)$, $\|y - x\| \leq \frac{\alpha}{1-\alpha}$. Hence, there exists a (δ_1, δ_2) -perturbation \mathcal{F}' of \mathcal{F} for which $R(x) \subseteq H'_q$ and $A'_q = A' + \frac{-\log(\gamma)}{\tau}I$. Since V is a Lyapunov function for the perturbed system, it holds that $V(\frac{1}{\gamma}e^{\tau A'}x) < V(x)$, that is, $V(x + \tau Ax) < \gamma V(x)$, concluding the proof. \square

PROOF of Proposition 13 Let V satisfy the conditions in Proposition 7. Without loss of generality, assume that $\mathcal{V}_{\max} = 1$. Hence, $\mathcal{V} \subseteq \mathbb{B}^*$. For each $x \in X$, let $c_x \in \mathcal{V}$ be such that $c_x^\top = V(x)$. Then, by (D1)–(D2), it holds that $\forall x \in X$, $c_x^\top x \geq \frac{1}{\epsilon}\|x\|$ and $V(x + \tau v) \leq \gamma c_x^\top x$. Since $\mathcal{V}_X \subseteq \mathcal{V}$ and $\frac{1}{\epsilon}\|x + \tau v\| \leq V(x + \tau v)$, it follows that $\forall x \in X$, $\forall v \in \mathcal{F}(x)$, $\forall c \in \mathcal{V}_X$, $c^\top(x + \tau v) \leq \gamma c_x^\top(x)$ and $\|x + \tau v\| \leq \gamma \epsilon c_x^\top x$. \square

PROOF of Proposition 19 Let $q \in Q$ and $x' \in X' \cap H_q$. Let $x \in X \cap H_q$ be such that $\|\hat{x} - \hat{x}'\| < r$, wherein $\hat{x} = \frac{x}{\|x\|}$ and $\hat{x}' = \frac{x'}{\|x'\|}$. Let $c \in \mathcal{V}(x')$. Let $v = A_q \hat{x}$ and $v' = A_q \hat{x}'$. We show that $c^\top v' < 0$. Note that $\|v - v'\| =$

$\|A_q(\hat{x} - \hat{x}')\| \leq \sigma r$, $|c^\top(\hat{x} - \hat{x}')| \leq r$ and $|V(\hat{x}) - V(\hat{x}')| \leq r$. Hence, $c^\top \tau \hat{x}' = c^\top(\hat{x}' + \tau \hat{x}') - V(\hat{x}') \leq c^\top(\hat{x} + \tau v) - V(x) + 2r + \tau\sigma r < c^\top(\hat{x} + \tau v) - V(x) + \frac{1-\gamma}{\epsilon} \leq 0$, where the before-last inequality comes from the assumption on r and the last inequality from the assumption on X and $\mathcal{V}_\circ \subseteq \frac{1}{\epsilon}\mathbb{B}^*$. Since x' , q and c were arbitrary, this concludes the proof. \square

PROOF of Theorem 21 Assume that Algorithm 1 produces at least $K+1$ counterexamples x_0, \dots, x_K . For each $k \in \{0, \dots, K\}$, let $q_k \in Q$ be such that $x_k \in H_{q_k}$ and $\min_{x \in X_k \cap H_{q_k}} \left\| \frac{x}{\|x\|} - \frac{x_k}{\|x_k\|} \right\| \geq r$ (by Corollary 20, such a q_k always exists). For each $q \in Q$, let $X_K \downarrow q = \{x_k : q_k = q\}$. By the pigeonhole principle, there is $q \in Q$ such that $|X_K \downarrow q| \geq (K+1)/|Q|$. Fix such a q . It holds that for all $x, y \in X_K \downarrow q$, if $x \neq y$ then $\|x - y\| \geq r$. Thus, $|X_K \downarrow q|$ is upper bounded by $\text{Pack}(r; \mathbb{S})$. This proves that $K+1 \leq |Q| \text{Pack}(r; \mathbb{S})$. \square