

Morse Graphs: Topological Tools for Analyzing the Global Dynamics of Robot Controllers

Ewerton R. Vieira^{2,3}, Edgar Granados¹, Aravind Sivaramakrishnan¹,
Marcio Gameiro², Konstantin Mischaikow², Kostas E. Bekris¹ **

¹ Dept. of Computer Science, Rutgers University, NJ, USA

² Dept. of Mathematics, Rutgers University, NJ, USA

³ DIMACS (Center for Discrete Mathematics and Theoretical Computer Science),
Rutgers University, NJ, USA

{er691,eg585,as2578,gameiro,mischaik,kostas.bekris}@rutgers.edu

Abstract. Understanding the global dynamics of a robot controller, such as identifying attractors and their regions of attraction (RoA), is important for safe deployment and synthesizing more effective hybrid controllers. This paper proposes a topological framework to analyze the global dynamics of robot controllers, even data-driven ones, in an effective and explainable way. It builds a combinatorial representation representing the underlying system’s state space and non-linear dynamics, which is summarized in a directed acyclic graph, the *Morse graph*. The approach only probes the dynamics locally by forward propagating short trajectories over a state-space discretization, which needs to be a Lipschitz-continuous function. The framework is evaluated given either numerical or data-driven controllers for classical robotic benchmarks. It is compared against established analytical and recent machine learning alternatives for estimating the RoAs of such controllers. It is shown to outperform them in accuracy and efficiency. It also provides deeper insights as it describes the global dynamics up to the discretization’s resolution. This allows to use the Morse graph to identify how to synthesize controllers to form improved hybrid solutions or how to identify the physical limitations of a robotic system.

Keywords: Topology, Robot Control, Robot Dynamics.

1 Introduction

Estimating Regions of Attraction (RoA) of a dynamical system is needed in robotics for understanding the conditions under which a controller can be safely applied to solve a task. It is also needed for composing controllers and forming hybrid solutions that work from a wider swath of the underlying state space. Estimating such RoAs, however, is challenging.

Computing a Lyapunov function (LF) can provide an RoA but obtaining an analytical expression for an LF is difficult for general non-linear systems. This motivated numerical solutions, which often still require access to the system’s

** The work is supported in part by an NSF HDR TRIPODS award 1934924.

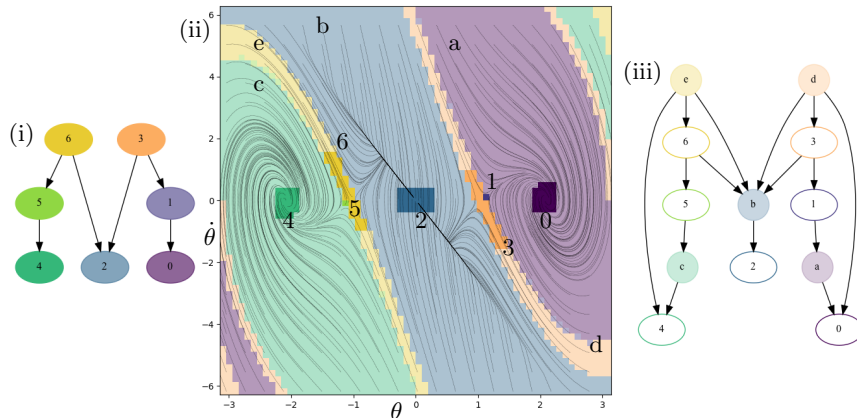


Fig. 1: Results of the proposed combinatorial analysis for a pendulum operating under an LQR controller so as to reach the $(0,0)$ state in the phase space. (i) The method outputs the Morse graph $MG(\mathcal{F})$ (left) to identify recurrent dynamics. (ii) The numbered Morse cells M represent recurrent sets associated with nodes of $MG(\mathcal{F})$. The RoAs of the regions in M are indicated by the corresponding light colors and letters. (iii) The graph on the right is used to identify the RoAs for the corresponding recurrent sets.

differential equation for computing an LF. Recent advances in data-driven control provide effective learned controllers [24, 21], which do not have analytical expressions. Machine learning methods have also been proposed to learn RoAs [13, 6] without access to the control law’s expression as well as for composing controllers [37, 7]. These methods, however, tend to be sensitive to parameters, are computationally demanding in time and memory while lacking guarantees. Therefore, there is a need for identifying RoAs, especially of data-driven systems, in a robust, mathematically rigorous, and computationally efficient way.

Contributions: This work sidesteps the estimation of an LF. It builds on top of recent progress in combinatorial dynamics and order theory [27–29] to propose a combinatorial analysis of the global dynamics of black-box robot controllers and describe attractors and their corresponding RoAs. The approach is based on a finite combinatorial representation of the state space and its non-linear dynamics. It only requires access to a discrete time representation of the dynamics and can handle analytical, data-driven and hybrid controllers.

For instance, consider the phase space of a pendulum as in Fig. 1 given an LQR controller for driving the system to the $(0,0)$ state. The proposed framework generates the combinatorial representation on the right, where nodes correspond to regions in a state space decomposition. This information is then summarized in more compact, annotated, acyclic directed graphs called Morse graphs, shown on the left. The nodes of Morse graphs can contain attractors of interest. The associated RoAs can also be inferred automatically from the combinatorial representation. The resulting information is finite and graphical in nature, thus, it can be easily queried and understood by a person. The accompanying evaluation shows that the proposed tools are relatively computationally efficient, provide a more global, explainable understanding of the dynamics, of-

ten achieve higher accuracy and provide stronger guarantees than alternatives. They also allow the composition of hybrid controllers with wider RoAs.

Related Work: Multiple numerical methods exist for estimating RoAs given direct access to the system’s expression [20]. For instance, maximal Lyapunov functions (LFs) [48] incrementally compute the RoA. Constructing an ellipsoidal RoA approximation reduces to a linear matrix inequalities (LMIs) problem [38, 39], which has been applied to wheeled robots [43] and NASA’s generic transport model [35]. There are also convex formulations that rely on LMI relaxations to solve a convex linear program and approximate the RoA of systems with polynomial dynamics and semi-algebraic inputs [25]. An LF can be restricted to be a sum-of-squares (SoS) polynomial constructed via semi-definite programming [36]. SoS methods can build randomized stabilized trees with LQR feedback [47], pre-compute funnel libraries [31] and acquire certificates of stability of rigid bodies with impacts and friction [41]. State space samples satisfying a Lyapunov-type inequality can construct neighborhoods where the candidate LF is certified [9]. In contrast, the proposed approach avoids computing an LF and does not require access to the analytical expression of the underlying control law.

Reachability analysis [5] computes the backward reachable tube of a dynamical system returning the maximal RoA without imposing shape restrictions. Tasks include RoA computation of dynamical walkers [14] and, combined with machine learning, maintain the system’s safety over a given horizon [22]. Barrier functions ensure safety of unknown dynamical systems, and can be learned with Gaussian Processes (GP) to obtain safe policies [3]. Barrier certificates (BCs) can identify areas needing exploration to expand the safe set [50].

Machine learning can be used to compute both LFs and BCs. One approach is to alternate between a learner and a verifier to search within a LFs set [12]. An alternative approximates the dynamics map as a piecewise linear neural network using a counterexample-guided method as a verifier to synthesize a LF [13]. LFs can also be constructed via stable data-driven Koopman operators [32]. LFs and BCs can be obtained by training a neural network and using an SMT solver as a verifier [1]. LFs for piecewise linear dynamical systems can be synthesized as the outputs of neural networks with leaky ReLU activations [17]. Given a system’s initial safe set, a neural network LF is trained to adapt to the RoA’s shape [44]. GPs can be used to obtain a Lyapunov-like function [30]. Finally, LFs can be synthesized while learning a controller to prove the controller’s stability and generate counter-examples to improve the controller [18]. This paper compares performance against a state-of-the-art ML approach that computes an LF [44].

This work builds on top of topological tools. Topology has been used for various problems in robotics, such as deformable manipulation [8, 4], robot perception [19], multi-robot problems [49], determining homotopy-inequivalent trajectories [40] and to extract higher-order dynamics for motion prediction [11]. Morse theory has been used to incrementally build local minima trees for multi-robot planning problems [34], and to find paths to cover 2D and 3D spaces [2]. To the best of the authors’ knowledge, this is the first application of recent advancements in topology to summarize the global dynamics of robot controllers.

2 Problem Setup

This work aims to systematically analyze the global dynamics of robot controllers based on combinatorial dynamics and order theory [27–29]. The prior theory is very general and applies to any continuous dynamical system defined over a locally compact metric space. The material is adopted and applied to the restricted setting of robot control problems. In particular, consider a nonlinear continuous-time system:

$$\dot{x} = f(x, u), \quad (1)$$

where $x(t) \in X$ is the state at time t in a domain $X \subseteq \mathbb{R}^n$, $u : X \rightarrow \mathbb{U} \subseteq \mathbb{R}^m$ is a Lipschitz continuous control as defined by a control policy $u(x)$, and $f : X \times \mathbb{U} \rightarrow \mathbb{R}^n$ is a Lipschitz continuous function, where \mathbb{U} is an open set in \mathbb{R}^m . The dynamical system consists of the model $f(\cdot)$, which can be accessed but it is not necessarily known analytically, and a control policy $u = u(x)$ that can be either analytical or learned from data.

For a given time $\tau > 0$, let $\phi_\tau : X \rightarrow X$ denote the function derived from solving Eq. (1) forward in time for duration τ from everywhere. Given that f and u are Lipschitz continuous, ϕ_τ is also Lipschitz continuous. Denote the global Lipschitz constant of ϕ_τ by L_τ . Observe that a RoA for Eq. (1) is a RoA under ϕ_τ . Therefore and w.l.o.g, the rest of this work focuses on the dynamics of ϕ_τ , which is not assumed, however, to be computable and available.

The objective is to identify a combinatorial approach, which can capture meaningful aspects of the dynamics of interest according to $\phi_\tau : X \rightarrow X$, which are continuous in nature. In this context, a subset of the state space $N \subset X$ is an attracting block for ϕ_τ , if $\phi_\tau(N) \subset \text{int}(N)$, where int denotes topological interior. This means that the system of Eq. (1) will not escape the subset N once it has entered it. Denote the set of attracting blocks of ϕ_τ by $\text{ABlock}(\phi_\tau)$. Given $N \in \text{ABlock}(\phi_\tau)$, its omega limit set is an invariant set for ϕ_τ defined as:

$$\omega(N) := \bigcap_{n \in \mathbb{Z}^+} \text{cl} \left(\bigcup_{k=n}^{\infty} \phi_\tau^k(N) \right)$$

where ϕ_τ^k is the composition $\phi_\tau \circ \dots \circ \phi_\tau$ (k times) and cl is topological closure. The attracting block N is a RoA for $\omega(N)$. In general $\text{ABlock}(\phi_\tau)$ is huge, containing uncountably many elements, and is too large to work with. Thus, the **problem** is to systematically identify a minimal finite subset of $\text{ABlock}(\phi_\tau)$ that both represents as tightly as possible the attractors and captures the maximal RoAs of these attractors.

Running Example: For exposition purposes, the following discussion will use the second-order pendulum as an example to explain the corresponding definitions and the proposed method (Fig 1). The pendulum is modeled by the differential equation $m\ell^2\ddot{\theta} = mG\ell \sin \theta - \beta\dot{\theta} + u$, with state $x := (\theta, \dot{\theta})$, where θ is the angle from the upright equilibrium $\theta_o = 0$, u is the input torque, m is the pendulum mass, G is the gravitational acceleration, ℓ is the pole length, and β is the friction coefficient. The control u in the running example is computed by the LQR approach described in Section 4. In Fig. 1, we use the time-1 map ϕ_1 of the flow of the pendulum under the LQR controller. The proposed method is not limited to this or similar low-dimensional systems/controllers.

3 Proposed Framework and Method

Overview: The method first approximates ϕ_τ by decomposing the state space X into regions ξ . For multiple initial states within each ξ the system is propagated forward for time τ to identify regions reachable from ξ . Given the reachability information and the Lipschitz continuity of ϕ_τ , a directed multi-valued graph representation \mathcal{F} stores each region ξ as a vertex and edges point from ξ to all regions in an outer (conservative) approximation of its true reachable set.

The method then computes the strongly connected components (SCC) of \mathcal{F} . An SCC is a maximal set of vertices of \mathcal{F} such that every pair of vertices in the SCC are reachable from each other. The non-trivial SCCs of \mathcal{F} , i.e., those with at least one edge, are called recurrent sets, and capture the recurrent dynamics of ϕ_τ . Every region ξ not in a recurrent set exhibits non-recurrent behavior. The same algorithm that computes SCCs also provides a topological sort of the vertices in \mathcal{F} , which allows to define reachability relationships between recurrent sets and non-recurrent regions. This gives rise to a condensation graph $\text{CG}(\mathcal{F})$, where all SCCs are condensed to a single vertex and edges reflect reachability according to the topological sort. Typically, $\text{CG}(\mathcal{F})$ is roughly the same size as \mathcal{F} and cumbersome to maintain. The implementation avoids explicitly storing either graph. The method succinctly captures the recurrent and non-recurrent dynamics in the Morse graph $\text{MG}(\mathcal{F})$, whose vertices are the recurrent sets of \mathcal{F} and whose edges reflect reachability according to the topological sort. Overall, the proposed method can be divided into the four steps described below:

- **Step 1.** State space decomposition and generation of input to represent ϕ_τ ;
- **Step 2.** Construction of the combinatorial representation \mathcal{F} of the dynamics given an outer approximation of ϕ_τ ;
- **Step 3.** Computation of Condensation Graph $\text{CG}(\mathcal{F})$ and Morse Graph $\text{MG}(\mathcal{F})$ via identification of recurrent sets/SCCs of \mathcal{F} and topological sort;
- **Step 4.** Derivation of RoAs for the recurrent sets from $\text{CG}(\mathcal{F})$;

Step 1 - State Space Decomposition and Generation of Input Data:

This paper considers the control system of Eq. (1) restricted to a state space given by an orthotope $X = \prod_{i=1}^n [a_i, b_i]$, allowing for the possibility of periodic boundary conditions. This allows for torus-like spaces, such as for the running example of the 2^{nd} -order pendulum. For simplicity, the accompanying implementation is using a uniform discretization of the state space based on 2^{k_i} subdivisions in the i -th component resulting in a decomposition of the state space into $\prod_{i=1}^n 2^{k_i}$ cubes of dimension n . The term \mathcal{X} denotes the collection of these cubes.

The method then generates the set of values of ϕ_τ at the corner points of cubes in \mathcal{X} . More precisely, let $V(\mathcal{X})$ denote the set of all corner points of cubes in \mathcal{X} . The method computes the set of ordered pairs $\Phi_\tau(\mathcal{X}) := \{(v, \phi_\tau(v)) \mid v \in V(\mathcal{X})\}$, by forward propagating the dynamics for time τ from all $V(\mathcal{X})$. In this way, this work does not assume exact, analytical knowledge of ϕ_τ , but rather exploits its existence. It only requires the ability to generate the set $\Phi_\tau(\mathcal{X})$.

Step 2 - Combinatorial Representation \mathcal{F} via outer approximation: The dynamics of the continuous ϕ_τ are approximated by a combinatorial multivalued

map $\mathcal{F}: \mathcal{X} \rightrightarrows \mathcal{X}$, where vertices are n -cubes $\xi \in \mathcal{X}$. The map \mathcal{F} contains directed edges $\xi \rightarrow \xi', \forall \xi' \in \Phi_\tau(\xi)$. The set of cubes identified by $\mathcal{F}(\xi)$ are meant to capture the possible states of $\phi_\tau(\xi)$. To obtain mathematically rigorous results about the dynamics of ϕ_τ , it is sufficient for \mathcal{F} to be an outer approximation of ϕ_τ , i.e., for

$$\phi_\tau(\xi) \subset \text{int}(\mathcal{F}(\xi)) \quad \text{for all } \xi \in \mathcal{X} \quad (2)$$

where int denotes topological interior. The left side of Eq. (2) indicates the set of states that can be achieved in time τ according to (1), which is unknown exactly. On the right side, $\mathcal{F}(\xi)$ is a list of vertices that can be identified with a set of n -cubes in X . The inclusion relation and int indicate the constraint that a large enough collection of cubes is chosen in $\mathcal{F}(\xi)$ to enclose $\phi_\tau(\xi)$ and at least an arbitrarily small overestimation is needed. The minimal outer approximation of ϕ_τ [26] is:

$$\mathcal{F}_{min}(\xi) := \{\xi' \in \mathcal{X} \mid \xi' \cap \phi_\tau(\xi) \neq \emptyset\},$$

that is, $\mathcal{F}_{min}(\xi)$ indicates the minimal set of cubes that contain the set of all states that can be reached in time τ starting in ξ . See Fig. 2(left) for an illustration. Even with complete knowledge of ϕ_τ , computation of \mathcal{F}_{min} is typically prohibitively expensive. Nevertheless, from a mathematical perspective it suffices to work with any $\mathcal{F}: \mathcal{X} \rightrightarrows \mathcal{X}$ that satisfies $\mathcal{F}_{min}(\xi) \subset \mathcal{F}(\xi)$ for all $\xi \in \mathcal{X}$. In general, the objective is to achieve a tight outer approximation, as the larger the size of the images of \mathcal{F} , the less the information about the dynamics of ϕ_τ .

The assumption for an outer approximation is that $\mathcal{F}(\xi) \neq \emptyset$ for all $\xi \in \mathcal{X}$. In practice it does not need to hold. Determining \mathcal{F} represents the major computational bottleneck as it involves numerical simulations of Eq. (1) or obtaining real-world experiments with the robotic system. The flexibility in the definition of an outer approximation provides flexibility in its construction. This work computes \mathcal{F} as follows: Recall that ϕ_τ




Fig. 2: The set $\phi_\tau(\xi)$ denotes the reachable states from all states in the cell ξ after time τ . Multivalued maps: (left) minimal/ideal outer approximation $\mathcal{F}_{min}(\xi)$ and (right) outer approximation obtained by using a Lipschitz constant.

is Lipschitz with constant $L = L_\tau$ (L depends on τ). Given $x \in X$, let $\overline{B}(x, \delta) = \{x' \in X \mid \|x - x'\| \leq \delta\}$ denote the δ -closed ball at state x . Define the diameter of $\xi \in \mathcal{X}$ by $d(\xi) := \max_{x, x' \in \xi} \|x - x'\|$ and the diameter of \mathcal{X} by $d := \max_{\xi \in \mathcal{X}} d(\xi)$. Note that for a uniform grid, $d = d(\xi)$, independently of the choice of ξ . Let $V(\xi)$ be the set of corner points of the cube ξ and:

$$\mathcal{F}(\xi) := \left\{ \xi' \mid \xi' \cap \overline{B}(\phi_\tau(v), Ld/2) \neq \emptyset \text{ for some } v \in V(\xi) \right\}.$$

Fig. 2(right) provides a relevant illustration. The definition of \mathcal{F} above is guaranteed to provide an outer approximation if L is (an upper bound for) the Lipschitz constant L_τ . In practice, however, only an estimate for the Lipschitz constant L_τ is available. In this case, an outer approximation can be obtained by evaluating ϕ_τ on a fine enough grid in ξ instead of just the corner points.

Step 3 - Identification of recurrent & non-recurrent behavior of \mathcal{F} :

Identifying all the recurrent sets \mathcal{M} of \mathcal{F} is performed using Tarjan's Strongly Connected Components (SCC) algorithm, which is linear in the number of elements of \mathcal{X} plus the number of edges in \mathcal{F} . The accompanying implementation uses a modified version of the algorithm, which does not store the whole digraph \mathcal{F} in memory and yet evaluates \mathcal{F} only once for each node [10, 33].

An indexing set \mathbf{P} is introduced in order to distinguish all the recurrent sets and helps to enumerate them: $\{\mathcal{M}(p) \mid p \in \mathbf{P}\}$. A partial order relation \leq on \mathbf{P} is imposed on the corresponding recurrent sets $\mathcal{M}(p)$. In particular, $q \leq p$ if there exists a path in \mathcal{F} from a $\xi \in \mathcal{M}(p)$ to a $\xi' \in \mathcal{M}(q)$. Identifying the partial order \leq for two recurrent sets is a question of reachability on \mathcal{F} between the two recurrent sets and is done taking advantage of the fact that Tarjan's algorithm also performs a topological sort on the vertices. For the examples of this paper, the number of recurrent sets is in the order of tens.

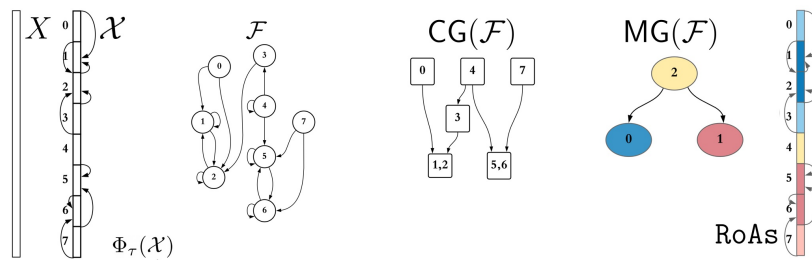


Fig. 3: From left to right: An example decomposition of a state space into regions and transitions between regions. The corresponding multivalued map \mathcal{F} . The condensation graph $\text{CG}(\mathcal{F})$ where SCCs have been condensed to a single vertex. The corresponding Morse graph where node 0 corresponds to the SCC $\{1,2\}$ with $\text{RoA}=\{0,1,2,3\}$ and node 1 corresponds to the SCC $\{5,6\}$ with $\text{RoA}=\{5,6,7\}$. Initial conditions in region 4 (identified with node 2) may end up either in regions $\{1,2\}$ (node 0) or $\{5,6\}$ (node 1).

Thus, the output of the SCC algorithm is a new graph representation, a condensation graph $\text{CG}(\mathcal{F})$ of \mathcal{F} , which is formed by contracting each strongly connected component of \mathcal{F} into a single vertex. The condensation graph $\text{CG}(\mathcal{F})$ is by definition a directed acyclic graph. The reachability on \mathcal{F} defines the direction of the edges in the condensation graph $\text{CG}(\mathcal{F})$ and relates to the partial order relation, where $v \leq w$, if there is a directed edge $w \rightarrow v$ in $\text{CG}(\mathcal{F})$. While the graph $\text{CG}(\mathcal{F})$ has condensed all SCCs into a single vertex, it is still a huge graph representation as it stores all the vertices of \mathcal{F} that are not in a recurrent set.

For this reason, the proposed method outputs the sub-graph derived only from the recurrent sets (i.e., the non-trivial SCCs) as these components are the only possible candidates for containing the attractors of RoAs given the level of discretization. This is referred to as the Morse graph $\text{MG}(\mathcal{F})$ of $\mathcal{F}: \mathcal{X} \rightrightarrows \mathcal{X}$ (shown in Fig. 1) and is the partially ordered set:

$$\text{MG}(\mathcal{F}) = \{\mathcal{M}(p) \subset \mathcal{X} \mid p \in (\mathbf{P}, \leq)\}. \quad (3)$$

Since (\mathbf{P}, \leq) is a partially ordered set, $\text{MG}(\mathcal{F})$ can be represented as a directed graph. The Morse graph $\text{MG}(\mathcal{F})$ of $\mathcal{F}: \mathcal{X} \rightrightarrows \mathcal{X}$ is the Hasse diagram of (\mathbf{P}, \leq) , i.e., the minimal directed graph from which (\mathbf{P}, \leq) can be reconstructed.

The Morse graph for the inverted pendulum is presented in Fig. 1 and is indexed by $\mathbf{P} = \{0, \dots, 6\}$ with order relations: $p < q$, iff there is a path from q to p in the digraph. Hence, 4, 2, and 0 are the minimal elements and $1 < 3$ and $5 < 6$. The $\text{MG}(\mathcal{F})$ is computed by Algorithm 1, which takes as input the decomposition \mathcal{X} , the dataset $\Phi_\tau(\mathcal{X})$ (representing the map ϕ_τ), and an estimate for the Lipschitz constant L of ϕ_τ . The dataset $\Phi_\tau(\mathcal{X})$ is used to compute the outer approximation \mathcal{F} , as described in Step 2. When the condensation graph is computed, the non-trivial SCCs (components with at least one edge) are flagged, as they become nodes of the Morse graph. Then, the only step remaining is to determine the reachability of recurrent sets as discussed above.

Algorithm 1: $\text{MorseGraph}(\mathcal{X}, \Phi_\tau(\mathcal{X}), L)$

```

1  $\mathcal{F} \leftarrow \text{OuterApproximation}(\mathcal{X}, \Phi_\tau(\mathcal{X}), L)$  //  $\mathcal{F}$  as a digraph but not
   stored in memory explicitly
2  $\text{SCC}(\mathcal{F}) \leftarrow \text{StronglyConnectedComponents}(\mathcal{F})$ 
3  $\text{CG}(\mathcal{F}) \leftarrow \text{CondensationGraph}(\text{SCC}(\mathcal{F}))$  // flag non-trivial SCCs
4  $\text{MG}(\mathcal{F}) \leftarrow \text{Reachability}(\text{CG}(\mathcal{F}))$ 
5 return  $\text{MG}(\mathcal{F}), \text{CG}(\mathcal{F})$ 
```

Step 4 - Derivation of RoAs: Define $O_\bullet: \mathcal{X} \rightrightarrows \mathbf{P}$ and $O^\bullet: \mathcal{X} \rightrightarrows \mathbf{P}$ as:
 $O_\bullet(\xi) := \min\{p \in \mathbf{P} \mid \text{there exists a path in } \mathcal{F} \text{ from } \xi \text{ to } \xi' \in \mathcal{M}(p)\}$ and
 $O^\bullet(\xi) := \max\{p \in \mathbf{P} \mid \text{there exists a path in } \mathcal{F} \text{ from } \xi \text{ to } \xi' \in \mathcal{M}(p)\}$.

Note that since \mathbf{P} is a poset it is possible that $O_\bullet(\xi)$ and $O^\bullet(\xi)$ have multiple values. Under the assumption that \mathcal{F} is an outer approximation of ϕ_τ , then, if $p \notin O_\bullet(\xi)$, it is true that for every $x \in \xi$ and any $n \geq 0$, $\phi_\tau^n(x) \notin \mathcal{M}(p)$.

Theorem 1. *If p is a minimal element of (\mathbf{P}, \leq) and $O^\bullet(\xi) = \{p\}$, then for every $x \in \xi$, there exists $n \geq 0$ such that $\phi_\tau^n(x) \in \mathcal{M}(p)$. As a consequence if p is a minimal element of (\mathbf{P}, \leq) , then $\{\xi \in \mathcal{X} \mid O^\bullet(\xi) = \{p\}\}$ is the maximal RoA for $\mathcal{M}(p)$ that can be rigorously identified using \mathcal{F} .*

Returning to the example of Fig. 1, $O^\bullet(\xi) = p$ for ξ is the region corresponding to the Morse set $\mathcal{M}(p)$, for $p = 0, \dots, 6$. For ξ in the RoA indicated by a, b, c, d , and e the following map arises: $O^\bullet(\xi) = 0$ for a , $O^\bullet(\xi) = 2$ for b , $O^\bullet(\xi) = 4$ for c , $O^\bullet(\xi) = 3$ for d , and $O^\bullet(\xi) = 6$ for e . This is indicated by the graph in Fig. 1(right). $O^\bullet(\xi)$ is the Morse node reachable from the corresponding region in the graph. It follows from Theorem 1 that the regions in Fig. 1 labeled 4 and c form the maximal RoA of 4, the regions labeled 2 and b form the maximal RoA of 2, and the regions labeled 0 and a form the maximal RoA of 0.

To obtain O^\bullet , the graph $\text{CG}(\mathcal{F})$ is explored with a depth first search (DFS) approach and for each visited vertex $v \in \text{CG}(\mathcal{F})$ the maximal reachable Morse nodes are identified. That is, the collection of $p \in \mathbf{P}$ such that there exists a path from v to a cube $\xi \in \mathcal{M}(p)$. See Appendix A for an implementation of the DFS.

Relationship to Continuous Dynamics: The condensation graph $\text{CG}(\mathcal{F})$ and the Morse graph $\text{MG}(\mathcal{F})$ of \mathcal{F} are highlighted in the right column of Fig. 4 as the combinatorial objects computed given the outer approximation \mathcal{F} of ϕ_τ . Each element of $\text{CG}(\mathcal{F})$ is identified with a region of X . This collection of regions is denoted by \mathbf{T} and, as shown in Fig 4, it is isomorphic as a poset to $\text{CG}(\mathcal{F})$. To find attracting blocks, the method uses the following fact [28]. Let \mathcal{F} be an outer

approximation of ϕ_τ : If $x \in T \in \mathbb{T}$ and $\phi_\tau(x) \in T' \in \mathbb{T}$, then $T' \leq T$ where \leq is the order relation on \mathbb{T} , i.e., forward orbits under ϕ_τ can be tracked by descending the order relation on \mathbb{T} . Stated more concisely: $\mathcal{O}(T) \in \text{ABlock}(\phi_\tau)$, where $\mathcal{O}(T) := \{T' \in \mathbb{T} \mid T' \leq T\}$ is the downset of $T \in \mathbb{T}$.

Dynamics **Combinatorics**




Fig. 4: Diagram relating the discretized dynamics in state space (left) to the proposed combinatorial representation of the global dynamics (right). The arrows \hookrightarrow represent inclusion maps and the arrows \cong indicate poset isomorphisms.

The set $\text{ABlock}(\phi_\tau)$ has the structure of a finite distributive lattice; that is, if $N, N' \in \text{ABlock}(\phi_\tau)$ then $N \cup N', N \cap N' \in \text{ABlock}(\phi_\tau)$. Furthermore, the collection $\{\mathcal{O}(T) \mid T \in \mathbb{T}\}$ generates a finite but huge sublattice of $\text{ABlock}(\phi_\tau)$. The elements of \mathbb{T} that are identified by elements of $\text{MG}(\mathcal{F})$ are denoted by \mathbb{M} and referred to as Morse cells (see Fig. 1). Given prior work [28], if \mathcal{F} is an outer approximation of ϕ_τ and $x \in X$ belongs to the chain recurrent set of ϕ_τ (recurrence allowing for an arbitrarily small error [15]), then x belongs to a Morse cell. Thus, the collection of recurrent sets of \mathcal{F} identifies the location in state space of recurrent dynamics of ϕ_τ . As in Fig. 4, \mathbb{M} inherits a partial order \leq from $\text{MG}(\mathcal{F})$. Its dynamical implications are derived from the dynamical implications of the partial order on \mathbb{T} . If $M, M' \in \mathbb{M}$, $M < M'$ and x is an initial condition that lies in M , then $\phi_\tau^n(x) \cap M' = \emptyset$ for all $n \geq 0$.

4 Results

The framework is compared against alternatives for estimating the RoAs, including numerical and machine learning methods that compute Lyapunov functions.

Robotic Systems: (i) A 2^{nd} -order (**Pendulum**), (ii) a 1^{st} -order car with (**Ackermann**) steering that cannot reverse [16], (iii) a 2^{nd} -order (**Acrobot**) [45]. The dynamics are simulated via numerical integration [23]. The state and control space limits are given in Table 1.

System	X	\mathbb{U}	Bounds on X	Bounds on \mathbb{U}	Controllers	Goal
Pendulum	$(\theta, \dot{\theta})$	\mathcal{T}	$[\{\pi, 2\pi\}, \{\pi, 2\pi\}]$	$[0.6372, 0.6372]$	Learned, LQR	$[0, 0]$
Ackermann	(x, y, θ)	(γ, V)	$[\{-10, -10, -\pi\}, \{10, 10, \pi\}]$	$[\{-\frac{\pi}{3}, 0\}, \{\frac{\pi}{3}, 30\}]$	Learned, LQR, Corke	$[0, 0, \frac{\pi}{2}]$
Acrobot	$(\theta_1, \theta_2, \dot{\theta}_1, \dot{\theta}_2)$	\mathcal{T}_2	$[\{0, -\pi, -6, -6\}, \{2\pi, \pi, 6, 6\}]$	$[14, 14]$	Hybrid, LQR	$[0, \pi, 0, 0]$

Table 1: Systems and controllers considered in the evaluation.

Controllers: For each system, alternative controllers are considered:

(i) An LQR controller linearizes the system around the goal: $\dot{x} = Ax + Bu$ ($A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{m \times n}$) and the controller $u = -Kx$ minimizes the cost $\mathcal{J}_{LQR} = x(T)^T F x(T) + \int_0^T (x(t)^T Q x(t) + u(t)^T R u(t)) dt$ ($F, Q \in \mathbb{R}^{n \times n}$, $R \in \mathbb{R}^{m \times m}$).

(ii) A Learned controller trained using the Soft Actor-Critic (SAC) [24] algorithm to maximize the expected return $\mathcal{J}(\pi) = \mathbb{E}_{\tau \sim \rho_\pi} [\sum_{t=0}^T \mathcal{R}(x_t)]$, where the reward function is $\mathcal{R} : X \rightarrow \{0, 1\}$. $\mathcal{R}(x_t) = 1$ iff x_t is within an ϵ distance from

the goal state and 0 otherwise.

(iii) A **Hybrid** controller for the **Acrobot** uses the **Learned** controller to drive the system to a relaxed goal distance, from where **LQR** takes over.

(iv) The **Corke** controller for **Ackermann** [16] transforms the state into polar coordinates (ρ, α, β) to define linear control laws for the velocity $v = K_\rho \rho$ and steering angle $\omega = K_\alpha \alpha + K_\beta \beta$. It’s able to drive a reverse-capable system to the goal assuming $K_\rho > 0, K_\beta < 0, K_\alpha - K_\rho > 0$. The experiments, however, consider only positive velocities limiting its reachability.

Comparison Methods: L-LQR and L-SoS use a linearized, unconstrained form of the dynamics to compute a Lyapunov function (LF) for the controller being considered. L-LQR uses the solution of the Lyapunov equation $v_{LQR}(x) = x^T P x$ for the linearized, unconstrained version of the system. L-SoS computes the LF according to $v_{sos}(x) = m(x)^T Q m(x)$ where $m(x)$ are monomials on x and Q is a positive semidefinite matrix. It uses SOSTOOLS [42] with SeDuMi [46] as the SDP solver. These methods cannot be used with data-driven controllers (like **Learned**) since they require a closed-form expression for the controller.

The Lyapunov Neural Network (L-NN) [44] is a state-of-the-art, machine learning approach with available software for identifying RoAs of black-box controllers. It returns a parametrized function that is trained to adapt to the RoA of a closed-loop dynamical system. Given an initial safe set around the desired equilibrium, a subset of non-safe states are forward propagated, classified and used to reshape the Lyapunov candidate in each iteration. The method needs access to a known safe-set and there is no guarantee the safe-region won’t shrink.

“Ground Truth” RoAs: For each benchmark (i.e., a controller-system pair), an approximation of the ground truth RoA for the goal state is computed by high-resolution discretization of the state space, and forward propagating the controller for a very long, fixed time horizon, or until the goal is reached. Appendix C provides the parameters for this ground truth evaluation.

Metrics: Given the “ground truth” RoA, the following metrics are reported:

(a) Table 2 provides the ratio of X ’s volume correctly identified to belong to the RoA (True Positives - TP) - its complement gives the ratio of X ’s volume *incorrectly* identified as not being in the RoA (False Negatives - FN);

(b) Table 3 provides the ratio of X ’s volume for which the dynamics have not been identified (Unidentified);

(c) Table 4 provides the amount of computational resources required for L-NN and **Morse Graph** (MG) measured using the number of forward propagations performed (steps) – the dominant computational primitive. Given their analytical nature, the L-LQR and L-SOS alternatives do not require forward propagations

Benchmark	L-NN	L-LQR	L-SOS	Ours: MG
Pend (LQR)	97.54%	69.91%	3.07%	97.49%
Pend (Learned)	30.18%			98.5%
Acro (LQR)	89.06%	26.84%	25.66%	96.36%
Acro (Hybrid)	13.79%			98.75%
Ack (LQR)	7.55%	21.78%	2.43%	0%
Ack (Corke)	10.23%		41.36%	86.69%
Ack (Learned)	91.47%			100%

Table 2: RoA ratios identified by the different methods. Best values per row in bold.

of the system and tend to be computationally faster but require access to an expression for the controller.

Quantitative Results: The proposed method tends to estimate larger volumes of the RoA compared to alternatives per Table 2. It also consistently identifies the dynamics for larger volumes of X compared to the comparison points, which cover lower volumes of X per Table 3. Moreover, the proposed method is broadly applicable to different controllers, and finds a larger volume of the RoA when compared to L-NN for the Pendulum (Learned), Acrobot (Hybrid), Ackermann (Learned).

Benchmark	L-NN	L-LQR	L-SOS	Ours: MG
Pend (LQR)	61.33%	61.33%	61.33%	1.66%
Pend (Learned)	81.44%			1.25%
Acro (LQR)	10.94%	73.16%	74.34%	3.64%
Acro (Hybrid)	86.21%			1.25%
Ack (LQR)	83.03%	82.87%	82.87%	100%
Ack (Corke)	27.81%		29.35%	24.5%
Ack (Learned)	8.53%			0.00%

Table 3: X 's ratio returned as *unidentified* by the different methods. Best values per row in bold.

Benchmark	L-NN	Ours: MG
Pend (LQR)	667.1M	6.6M
Pend (Learned)	341.9M	6.6M
Acro (LQR)	5.7B	1.1B
Acro (Hybrid)	533M	2.1B
Ack (LQR)	9.9M	520M
Ack (Corke)	37.5M	13M
Ack (Learned)	704.6M	520M

Table 4: Number of propagation required. Best values per row in bold.

The computational needs of the Morse Graph is one or two orders of magnitude less than that of the L-NN. The learned controllers benefit the most from the topological approach as it provides in all cases a higher coverage of the RoA with fewer propagations. There are two cases where MG takes a larger number of propagation steps compared to L-NN. In the case of Ackermann (LQR), this is because MG is unable to find a unique attractor for the system (see discussion below). In the case of Acrobot (Hybrid), L-NN fails to identify the true RoA accurately.

Some of the comparison points may incorrectly identify a volume of X as belonging to RoA (False Positives - FP). This is not true for the Pendulum, since the attractor of interest is not at the boundary of the RoA. But for the LQR and Corke controllers of the Ackermann, the desired goal is not an attractor since some trajectories close to the goal region escape X (a consequence of not allowing negative velocities). Therefore, the comparison methods fail to conservatively estimate the RoA, resulting in FPs. The topological framework, however, does not result in FPs (explained below). The Ackermann Learned controller and the Acrobot controllers present no FPs since the goal region is an attractor.




Fig. 5: RoA estimated by L-NN (1st and 3th image) and Morse Graph (2nd and 4th image, where the dark shading is the attractor) for the LQR (left) and Learned controllers (right) of the pendulum.

Pendulum Study: The RoAs compute controllers for the pendulum are shown in Fig 5. For **Morse Graph**, the attractor discovered is shown at the center of the state space, and it contains the goal region. Note that **L-NN** does not find the attractor, but assumes one exists containing the goal region. Moreover, **L-NN** is unable to cover the full RoA for the Learned controller, making it an undesirable solution. **Morse Graph** is also shown to work well when (a subset of) the state space is periodic. For instance, the RoA of the Pendulum (LQR) also includes the regions at the corners of the planar representation of the cylinders in Fig 5 (Left). This is captured by **MorseGraph**, but not by **L-NN**.

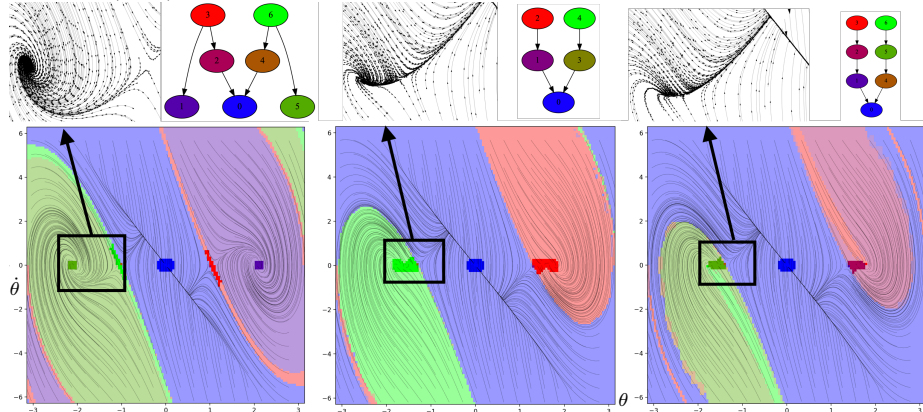


Fig. 6: RoA for the Pendulum when the upper torque bound is: (left) 0.637; (center) 0.724; (right) 0.736. For all torques the RoA found is 97% of the true one.

Pendulum with different torques: Fig 6 illustrates the robustness of the topological framework to different bounds for the allowed torque of the LQR controller. The proposed method consistently finds the RoA and the attractor that contains the goal region. Moreover, it persistently covers 99% of the true RoA even when the dynamics change, for instance in Fig. 6 one attractor and one saddle eventually collide and disappear.

Ackermann Study:

Learned controller for Ackermann: Although the RoA is the full state space, only the Morse Graph identifies this. The unique attractor obtained (right) also provides insight into how the controller works. For a small resolution of the discretization of X , it presents as a torus-like shape, which suggests recurrent behavior. This shape can be explained by the behavior of the car when it gets close to the goal region with the wrong orientation, where it tries to fix this orientation by performing a loop. For a more refined discretization, the method can distinguish long trajectories from recurrent behavior, finding a smaller attractor that contains the goal region.




Fig. 7: For Ackermann (Learned), the entire state space is correctly identified by the Morse Graph as an RoA for the shown Unique Morse set (Torus-like shape).

LQR controller for Ackermann: The Morse Graph results in a large, unique attractor that contains 74% of X as the discretization is proven insufficient. The comparison points fail by producing false positives (FP). The proposed method is conservative and safe. It avoids FPs but needs more subdivisions for a more comprehensive understanding of these global dynamics.

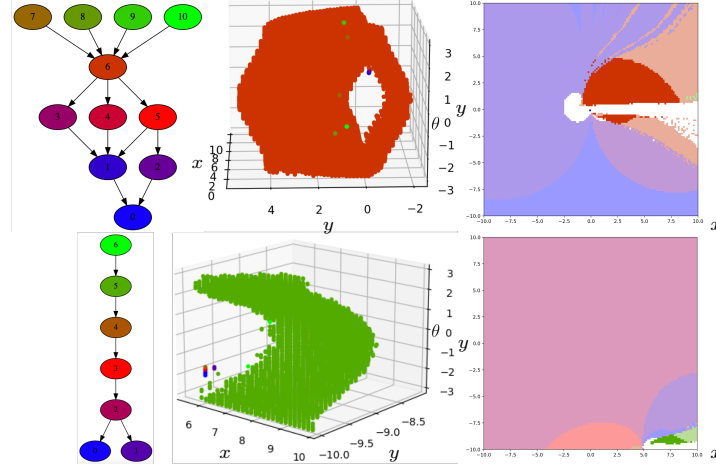


Fig. 8: Morse Graph, Morse sets and 2D projection on $\theta = 0$ of the RoA for the Corke controller applied to the Ackermann. (top row) Corke controller with the goal is set to be $(0, 0, 1.57)$ and (bottom row) Corke controller with the goal set to be $(6, -10, -1.57)$.

Corke controller for Ackermann: Some trajectories leave X in this experiment and modifications are needed to compute the RoA. A node \star is included in $\text{CG}(\mathcal{F})$, and for every cube $\xi \in \mathcal{X}$ s.t. $\Phi_\tau(\xi) \cap X^c \neq \emptyset$, edges are added from ξ to \star before computing the Morse Graph. Thus, \star is a minimal node of $\text{MG}(\mathcal{F})$. Then, a modified RoA computation is applied where O^\bullet and \max are changed to O_\bullet and \min , respectively, and the resulting output is O_\bullet . Finally, for each element $R \in \text{mRoA}$, the method computes $R_\bullet = R - R_\star$, where $R_\star = \{\xi \in \mathcal{X} \mid O_\bullet(\xi) = \{\star\}\}$. Consequently, for each maximal RoA, the cubes in R_\star are removed since they have some trajectories that escape X . So, R_\bullet is a conservative estimate of the RoA. In Fig 8, the white region is the set of cubes in R_\bullet .

Devising a Hybrid controller for Ackermann given the Morse Graph output: Given the information from the Morse Graph, it is possible to synthesize a hybrid controller that has a bigger RoA than the original $\text{ROA}_{\text{Corke}}$ of the Corke Controller u_{init} . The strategy selects a state in $\text{ROA}_{\text{Corke}}$, different than the original goal, as the goal for a new Corke controller u_{inter} . Define as $\text{RoA}_{\text{inter}}$ the RoA of the new Corke controller. If $\text{RoA}_{\text{inter}}$ overlaps with $X - \text{RoA}_{\text{Corke}}$, the hybrid controller is defined as: for states in $\text{ROA}_{\text{Corke}}$, apply u_{init} , and for states in $X - \text{RoA}_{\text{Corke}}$, apply u_{inter} , and then apply u_{init} when the system enters $\text{ROA}_{\text{Corke}}$. Fig 8 (bottom) shows the RoA for u_{inter} with goal $(6, -10, -1.57)$, which is in the $\text{ROA}_{\text{Corke}}$. A new controller u_{inter} was devised for this goal and its $\text{RoA}_{\text{inter}}$ contained $X - \text{RoA}_{\text{Corke}}$. The integration of u_{inter} with the original u_{init} result in a hybrid solution that covers the entire state space, which was verified empirically using the Morse Graph.

Acrobot study: The Morse graph for both controllers of the Acrobot is shown in Fig 9. They have a high success rate: the estimated RoA for node 0 covers more than 93% of the true RoA. The possible recurrent dynamics described by node 1 are long trajectories interpreted as recurrent by the proposed method. This can be addressed by increasing the state space discretization for an additional computational cost.



Fig. 9

Hybrid Controller for Acrobot: The Acrobot learned controller is unable to find a solution within a given time horizon for a goal condition: $\mathcal{B}(x_G, 0.1)$. Hence, a hybrid solution is proposed. The learned controller is first applied until the system reaches a relaxed goal: $\mathcal{B}(x_G, 0.6)$. The proposed method finds that the RoA for the relaxed goal condition is 100% of the state space, and is inside the RoA of the LQR controller. Hence, once the trajectory reaches $\mathcal{B}(x_G, 0.6)$ using the learned controller, LQR is applied to get into the smaller region $\mathcal{B}(x_G, 0.1)$. This results in a hybrid controller with a 100% success rate and a reduction in trajectory length relatively to just executing LQR of around 50%.

5 Conclusion

This work presents a novel method based on topology to identify attractors and their RoAs for robotic systems controlled by black-box controllers. Experimental evaluation on simulated benchmarks shows that the proposed method efficiently identifies the global dynamics with fewer samples from the dynamics model compared to data-driven alternatives. It does not require knowledge of the system or controller dynamics, such as differentiability or the guaranteed presence of an attractor for the system's goal region. This makes it suitable for data-driven controllers, where it significantly outperforms alternatives in identifying their RoA. Moreover, the proposed method provides a compact description of the global dynamics, which allows to compose multiple controllers into hybrid solutions that reach the goal from the full state space. The evaluation section presents two such hybrid controllers designed based on the Morse Graph output that yield notable properties: one increased the RoA to the whole state space; and the other decreased the length of solution trajectories by half.

Even though the proposed method requires less number of samples from the dynamics model, it still explores the entirety of the state space. This can potentially be mitigating for large and high-dimensional state spaces. Future work will explore extensions of the current topological approach, where the focus will be on finding the RoA for a single attractor, thereby requiring less computational resources. Integration with Gaussian Processes and machine learning primitives can help identify a smaller set of states where the system is propagated from so as to further reduce data requirements.

References

1. Abate, A., Ahmed, D., Edwards, A., Giacobbe, M., Peruffo, A.: Fossil: a software tool for the formal synthesis of lyapunov functions and barrier certificates using neural networks. In: HSCC. pp. 1–11 (2021)

2. Acar, E.U., Choset, H., Rizzi, A.A., Atkar, P.N., Hull, D.: Morse decompositions for coverage tasks. *IJRR* **21**(4), 331–344 (2002)
3. Akametalu, A.K., Fisac, J.F., Gillula, J.H., Kaynama, S., Zeilinger, M.N., Tomlin, C.J.: Reachability-based safe learning with gaussian processes. In: CDC (2014)
4. Antonova, R., Varava, A., Shi, P., Carvalho, J.F., Kragic, D.: Sequential topological representations for predictive models of deformable objects. In: L4DC (2021)
5. Bansal, S., Chen, M., Herbert, S., Tomlin, C.J.: Hamilton-jacobi reachability: A brief overview and recent advances. In: CDC (2017)
6. Berkenkamp, F., Moriconi, R., Schoellig, A.P., Krause, A.: Safe learning of RoAs for uncertain, nonlinear systems with Gaussian Processes. In: CDC (2016)
7. Berkenkamp, F., Schoellig, A.P.: Safe and robust learning control with Gaussian Processes. In: ECC (2015)
8. Bhattacharya, S., Kim, S., Heidarsson, H., Sukhatme, G.S., Kumar, V.: A topological approach to using cables to manipulate sets of objects. *IJRR* **34**(6) (2015)
9. Bobiti, R., Lazar, M.: A sampling approach to constructing lyapunov functions for nonlinear continuous-time systems. In: CDC (2016)
10. Bush, J., Gameiro, M., Harker, S., Kokubu, H., Mischaikow, K., Obayashi, I., Pilarczyk, P.: Combinatorial-topological framework for the analysis of global dynamics. *Chaos: An Interdisciplinary Journal of Nonlinear Science* **22**(4) (2012)
11. Carvalho, J.F., Vejdemo-Johansson, M., Pokorný, F.T., Kragic, D.: Long-term prediction of motion trajectories using path homology clusters. In: IROS (2019)
12. Chen, S., Fazlyab, M., Morari, M., Pappas, G.J., Preciado, V.M.: Learning lyapunov functions for hybrid systems. In: HSCC. pp. 1–11 (2021)
13. Chen, S., Fazlyab, M., Morari, M., Pappas, G.J., Preciado, V.M.: Learning region of attraction for nonlinear systems. arXiv preprint arXiv:2110.00731 (2021)
14. Choi, J.J., Agrawal, A., Sreenath, K., Tomlin, C.J., Bansal, S.: Computation of RoAs for Hybrid Limit Cycles Using Reachability. arXiv:2201.08538 (2022)
15. Conley, C.: Isolated invariant sets and the Morse index, CBMS Regional Conference Series in Mathematics, vol. 38. American Math. Society, R.I. (1978)
16. Corke, P.I., Khatib, O.: Robotics, vision and control: fundamental algorithms in MATLAB, vol. 73. Springer (2011)
17. Dai, H., Landry, B., Pavone, M., Tedrake, R.: Counter-example guided synthesis of neural network lyapunov functions for piecewise linear systems. In: CDC (2020)
18. Dai, H., Landry, B., Yang, L., Pavone, M., Tedrake, R.: Lyapunov-stable neural-network control. arXiv preprint arXiv:2109.14152 (2021)
19. Ge, Q., Richmond, T., Zhong, B., Marchitto, T.M., Lobaton, E.J.: Enhancing the morphological segmentation of microscopic fossils through localized topology-aware edge detection. *Autonomous Robots* **45**(5), 709–723 (2021)
20. Giesl, P., Hafstein, S.: Review on computational methods for lyapunov functions. *Discrete & Continuous Dynamical Systems-B* **20**(8), 2291 (2015)
21. Gillen, S., Molnar, M., Byl, K.: Combining deep reinforcement learning and local control for the acrobot swing-up and balance task. In: CDC (2020)
22. Gillulay, J.H., Tomlin, C.J.: Guaranteed safe online learning of a bounded system. In: IROS (2011)
23. Granados, E., Sivaramakrishnan, A., McMahon, T., Littlefield, Z., Bekris, K.E.: Machine learning for kinodynamic planning (ml4kp) software (2021)
24. Haarnoja, T., Zhou, A., Abbeel, P., Levine, S.: Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In: ICML (2018)
25. Henrion, D., Korda, M.: Convex computation of the RoA of polynomial control systems. *IEEE Tran. on Automatic Control* **59**(2), 297–312 (2013)

26. Kalies, W.D., Mischaikow, K., Vandervorst, R.: An algorithmic approach to chain recurrence. *Found. Comput. Math.* **5**(4), 409–449 (2005)
27. Kalies, W.D., Mischaikow, K., Vandervorst, R.: Lattice structures for attractors I. *J. Comput. Dyn.* **1**(2), 307–338 (2014)
28. Kalies, W.D., Mischaikow, K., Vandervorst, R.: Lattice structures for attractors II. *Foundations of Computational Mathematics* **1**(2), 1–41 (2015)
29. Kalies, W.D., Mischaikow, K., Vandervorst, R.: Lattice structures for attractors III. *Journal of Dynamics and Differential Equations* pp. 1572–9222 (2021)
30. Lederer, A., Hirche, S.: Local Asymptotic Stability Analysis and Region of Attraction Estimation with Gaussian Processes. In: *CDC* (2019)
31. Majumdar, A., Tedrake, R.: Funnel libraries for real-time robust feedback motion planning. *The International Journal of Robotics Research* **36**(8), 947–982 (2017)
32. Mamakoukas, G., Abraham, I., Murphey, T.D.: Learning stable models for prediction and control. *IEEE Trans Robot* (2020)
33. Marcio Gameiro, S.H.: CMGDB: Conley Morse Graph Database Software (2022)
34. Orthey, A., Toussaint, M.: Visualizing local minima in multi-robot motion planning using multilevel morse theory. In: *International Workshop on the Algorithmic Foundations of Robotics*. pp. 502–517. Springer (2020)
35. Pandita, R., Chakraborty, A., Seiler, P., Balas, G.: Reachability and RoA analysis applied to GTM dynamic flight envelope assessment. In: *AIAA CNC* (2009)
36. Parrilo, P.A.: Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization. *California Institute of Technology* (2000)
37. Perkins, T.J., Barto, A.G.: Lyapunov design for safe reinforcement learning. *JMLR* **3**(Dec), 803–832 (2002)
38. Pesterev, A.V.: Attraction domain estimate for single-input affine systems with constrained control. *Automation and Remote Control* **78**(4), 581–594 (2017)
39. Pesterev, A.V.: Attraction domain for affine systems with constrained vector control closed by linearized feedback. *Autom. & Remote Control* **80**(5) (2019)
40. Pokorný, F.T., Kragic, D., Kavraki, L.E., Goldberg, K.: High-dimensional winding-augmented motion planning with 2D topological task projections & persistent homology. In: *ICRA* (2016)
41. Posa, M., Tobenkin, M., Tedrake, R.: Lyapunov analysis of rigid body systems with impacts and friction via sums-of-squares. In: *HSCC*. pp. 63–72 (2013)
42. Prajna, S., Papachristodoulou, A., Parrilo, P.A.: Introducing SOSTOOLS: A general purpose sum of squares programming solver. In: *CDC* (2002)
43. Rapoport, L.B., Morozov, Y.V.: Estimation of attraction domains in wheeled robot control using absolute stability approach. *IFAC* **41**(2), 5903–5908 (2008)
44. Richards, S.M., Berkenkamp, F., Krause, A.: Lyapunov Neural Network: Adaptive stability certification for safe learning of dynamical systems. In: *CoRL* (2018)
45. Spong, M.: The swing up control problem for the acrobat. *IEEE Control Systems Magazine* **15**(1), 49–55 (1995). <https://doi.org/10.1109/37.341864>
46. Sturm, J.F.: Using sedumi 1.02, a matlab toolbox for optimization over symmetric cones. *Optimization methods and software* **11**(1-4), 625–653 (1999)
47. Tedrake, R., Manchester, I.R., Tobenkin, M., Roberts, J.W.: LQR-trees: Feedback motion planning via sums-of-squares verification. *IJRR* **29**(8) (2010)
48. Vannelli, A., Vidyasagar, M.: Maximal lyapunov functions and domains of attraction for autonomous nonlinear systems. *Automatica* **21**(1), 69–80 (1985)
49. Varava, A., Hang, K., Kragic, D., Pokorný, F.T.: Herding by caging: a topological approach towards guiding moving agents via mobile robots. In: *R:SS* (2017)
50. Wang, L., Theodorou, E.A., Egerstedt, M.: Safe learning of quadrotor dynamics using barrier certificates. In: *ICRA* (2018)