On the Limits of Sampling-Based Reachability: Geometry, Dynamics, and Sample Complexity

Anonymous Author(s)
Affiliation
Address
email

Abstract: Reachability analysis is central to safety-critical control, robotics, and neural network verification, but classical computational methods, such as Hamilton–Jacobi reachability and set propagation, scale poorly with state dimension. Sampling-based methods have emerged as a promising alternative, often providing finite-sample guarantees that bound the probability-mass left uncovered. However, an explicit account of how the geometry of the initial set, the dynamics, and the sampling law affect the accuracy of the estimator is not fully available in the literature. We study this by casting sampling-based reachable-set recovery as geometric support estimation over a family of problems specified by an initial set, its dynamics, and a sampling law. First, we identify two regularity properties, positive reach of the initial set's complement and Lipschitz continuity of the dynamics, that together make recovery well-posed: a probability-mass coverage guarantee can be upgraded to accuracy r in Hausdorff distance. Second, we bound the resulting sample complexity: recovery is achievable with $\tilde{\mathcal{O}}\big((e^{3LT}/r)^n\big)$ samples, exponential in both the state dimension and the time horizon. Third, we show that neither can be removed: an almost-matching minimax lower bound of $\Omega\big((e^{LT}/r)^n\big)$ holds for every estimator, so the exponential dependence on dimension and the degradation over the horizon are both intrinsic, not artifacts of a particular method. Experiments on nonlinear systems confirm that adversarial sampling improves constants but not the scaling.

Keywords: Reachability analysis, sample complexity, positive reach

1 Introduction

Reachability analysis is a fundamental tool for certifying the safety of dynamical and learning-enabled systems [1, 2, 3, 4]. Classical computational methods, based on Hamilton–Jacobi reachability [5, 6] or set propagation [7, 8], give strong deterministic guarantees but scale poorly with state dimension. This barrier is especially pronounced in modern learning-enabled systems, where neural-network controllers and learned dynamics put classical methods out of reach. To improve scalability, sampling-based methods approximate reachable sets by propagating finitely many initial states through the dynamics [9, 10, 11, 12]. They are model-agnostic, parallelizable, and easy to implement, applicable to general nonlinear and learning-based systems. These methods are part of a broader rise of sampling-based control and optimization for high-dimensional systems [13, 14, 15], made practical by modern parallel hardware.

The accompanying finite-sample guarantees, however, are limited. They either bound the probability-mass the estimate fails to cover [16, 17, 18], permitting an estimate to miss a thin, spatially separated region without violating the bound, or they measure error in Hausdorff distance against the convex hull of the reachable set [19, 20], which over-approximates nonconvex sets by

All codes and relevant materials are available at: https://github.com/AnonyDreamer/reachapprox

Submitted to the 10th Conference on Robot Learning (CoRL 2026). Do not distribute.

filling holes and merging disconnected components. What remains missing is a characterization of Hausdorff recovery of the actual set, and an account of how the sample cost depends on the geometry of the initial set, the dynamics, and the sampling law. Posed this way, the question becomes one of geometric support estimation under flow distortion, a perspective with a rich classical theory [21, 22, 23] that has not yet been brought to bear on the reachability setting.

Contributions. To fill this gap, we study these questions for autonomous systems $\dot{x} = F(x)$, recovering the reachable set $S_T = \mathcal{R}_T(S_0)$ from endpoint samples of trajectories initialized in S_0. Casting recovery as geometric support estimation over a family of problems specified by the initial set, its dynamics, and a sampling law, we make the following contributions:

1. **Well-posedness.** We identify two regularity properties, positive reach of the initial set's complement (r_0) and Lipschitz continuity of the dynamics (L), under which a probability-mass coverage guarantee can be upgraded to Hausdorff-accurate recovery of the reachable set.
2. **Sample complexity.** We give a random-covering upper bound showing that $\tilde{\mathcal{O}}\big((e^{3LT}/r)^n\big)$ endpoint samples suffice for r-accurate inner coverage of the reachable set. Together with an estimator-dependent outer-deviation condition, this yields a Hausdorff guarantee.
3. **Unavoidability.** We prove an almost-matching minimax lower bound: any estimator, under any sampling law, requires $\Omega\big((e^{LT}/r)^n\big)$ samples on some instance, so the dimension dependence and the horizon degradation are intrinsic, not artifacts of a method.
4. **Numerical experiments.** On a 2D non-Lipschitz example outside our assumptions and a multi-link robot arm in MuJoCo, the empirical sample complexity tracks the predicted dimension-dependent scaling, and adversarial sampling improves the constants but not the scaling.

Closely related work. Our analysis is closely related to that of Lew et al. [20], who derive a finite-sample Hausdorff guarantee for a sampling-based estimator of the convex hull of the reachable set under essentially the same problem family as ours: the complement of the initial set is r-convex, the vector field is Lipschitz, and the sampling law admits a Lebesgue density lower bound. Our positive-reach assumption $\mathrm{reach}(S_0^c) \geq r_0$ (Definition 1) is strictly stronger and implies their r-convexity. We go further in three directions: (i) we study Hausdorff recovery of the actual endpoint support, rather than the convex hull of the reachable set; (ii) we make the horizon dependence explicit, rather than absorbing it into a single Lipschitz constant of the flow; and (iii) we certify the rate with a matching minimax lower bound that holds for every sampling law. The r^{-n} rate is classical in geometric support estimation [21, 22, 23] and is known to be minimax-optimal under positive-reach conditions [24]; in our setting it persists under flow distortion with an explicit e^{nLT} amplification. A more extensive discussion of adjacent literatures, including scenario optimization, conformal prediction, and Lipschitz-based bloating, is provided in Appendix A.

Notation. For a Lebesgue measurable set $A \subseteq \mathbb{R}^n$, we write $A^c := \mathbb{R}^n \setminus A$ for its complement, $\overline{A}$ for its closure, and ∂A for its boundary. Let $|A| := \lambda_n(A)$, where λ_n is the n-dimensional Lebesgue measure. For $x \in \mathbb{R}^n$, $r > 0$, and nonempty sets $A, B \subseteq \mathbb{R}^n$, define $\mathcal{B}_r(x) := \{y \in \mathbb{R}^n : \|y - x\| \leq r\}$, $d(x, A) := \inf_{a \in A} \|x - a\|$, and $\mathcal{B}_r(A) := \{x \in \mathbb{R}^n : d(x, A) \leq r\}$. The Hausdorff distance between A and B is $d_H(A, B) := \max\{\sup_{a \in A} d(a, B),\ \sup_{b \in B} d(b, A)\}$. Let $\omega_n := |\mathcal{B}_1(0)|$ denote the volume of the unit ball in $\mathbb{R}^n$. We use $A \oplus B := \{a + b : a \in A,\ b \in B\}$, $A \ominus B := \{x \in \mathbb{R}^n : x + B \subseteq A\}$, and $A \circ B := (A \ominus B) \oplus B$ for Minkowski sum, erosion, and opening, respectively. For a probability measure P, $\mathrm{supp}(P)$ denotes its support, and for probability measures P, Q, $D_{\mathrm{KL}}(P\|Q)$ denotes the Kullback–Leibler divergence.

2 Problem Formulation

We consider an autonomous continuous-time dynamical system $\dot{x} = F(x)$, where $x \in \mathcal{X} \subseteq \mathbb{R}^n$ and $F : \mathcal{X} \to \mathbb{R}^n$. Given an initial condition $x_0 \in \mathcal{X}$, let $\varphi(t, x_0)$ denote the solution at time $t \geq 0$. For an initial set $S_0 \subseteq \mathcal{X}$ and a time horizon $T > 0$, the forward reachable set is

$$\mathcal{R}_T(S_0) := \{\varphi(T, x_0) \in \mathcal{X} :\ x_0 \in S_0\}. \tag{1}$$

When S_0 and T are known, we write $S_T := \mathcal{R}_T(S_0)$. This paper studies the problem of approximating S_T from finitely many endpoint samples. Specifically, suppose that initial states $X_1, \ldots, X_N$ are sampled independently from a probability distribution P_0 supported on $\overline{S_0}$, and the corresponding endpoints are $Y_i := \varphi(T, X_i), i = 1, \ldots, N$. These endpoints are i.i.d. samples from the push-forward distribution $P_T := (\varphi(T, \cdot))_{\#} P_0$, whose support is $\overline{S_T}$. A sampling-based reachable-set estimator is a measurable set-valued map $\widehat{S}_N = \widehat{S}_N(Y_1, \ldots, Y_N) \subseteq \mathbb{R}^n$ that approximates the reachable set S_T.

Many existing sampling-based and learning-based approaches provide probabilistic coverage guarantees. Informally, such results show that, with probability at least $1 - \delta$, the estimated set misses at most ε probability-mass under the endpoint distribution P_T. A typical guarantee takes the form $P_T(\overline{S_T} \setminus \widehat{S}_N) \leq \varepsilon$ with probability larger than $1 - \delta$ over the sampled endpoints and the joint distribution P_T^N of $(Y_1, \ldots, Y_N)$ [16, 18]. However, probability-mass accuracy alone does not imply geometric accuracy. An estimator may miss only a small amount of probability-mass while still incurring a large Hausdorff error. This can happen when the missed region has small probability but is geometrically far from the estimated set, or when the reachable set contains thin cusps, spikes, or low-density regions. In safety and verification problems, such geometric errors are important: missing a small-probability but spatially significant region may still lead to an incorrect reachable-set certificate.

Motivated by this gap, our goal is to understand how many endpoint samples are necessary and sufficient for $\widehat{S}_N$ to approximate S_T in Hausdorff distance. To be more specific, for a target accuracy $r > 0$ and confidence level $1 - \delta$, we seek conditions and sample size N under which:

$$P_T^N \left(d_H(S_T, \widehat{S}_N) \leq r \right) \geq 1 - \delta. \tag{2}$$

This objective is stronger than controlling the probability-mass of the missed region. It requires every point of the true reachable set to lie within distance r of the estimator, and conversely requires the estimator not to extend too far away from the true reachable set. This property is arguably necessary for safety-critical circumstances.

The central question is therefore: when can a probability-mass guarantee be converted into a Hausdorff-distance guarantee? Such a conversion is impossible without additional structure. Indeed, if there exists a point $x \in S_T$ such that $\mathcal{B}_r(x) \cap \overline{S_T}$ has arbitrarily small P_T-mass, then an estimator may miss this entire r-scale neighborhood while still satisfying a small probability error bound. In that case, the probability error is small, but the Hausdorff distance from S_T to $\widehat{S}_N$ is at least of order r. Thus, to make (ε, δ)-type probability guarantees geometrically meaningful, one needs a lower bound on the probability-mass of local geometric neighborhoods.

Figure 1 illustrates this phenomenon on the autonomous system $\dot{y} = 0, \dot{x} = x^2$. Starting from two initial sets with comparable geometric scale, a disk and a nonconvex star, we uniformly sample initial conditions, propagate the samples, and construct support estimators from the endpoints. The flow map expands regions with large x, causing geometric error to increase with time, i.e. the Hausdorff distance grows super-exponentially as the time horizon increases, reflecting the strong expansion of the flow map. Moreover, for the same sampling budget and comparable geometric scale, the star-shaped set produces larger Hausdorff error than the disk. This behavior is not explained by probability-mass alone. Rather, it reflects

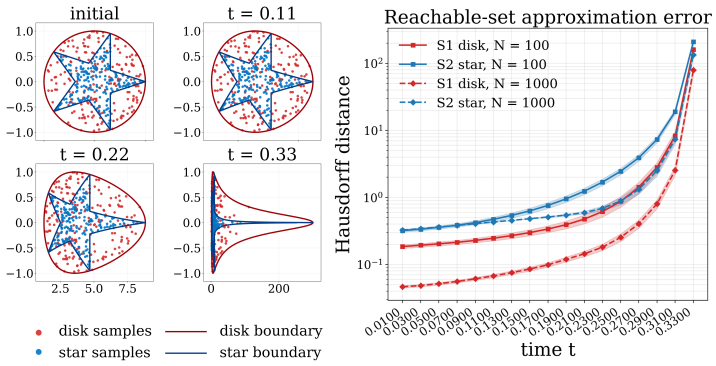

Figure 1: Reachable-set approximation under the autonomous dynamics $\dot{y} = 0, \dot{x} = x^2$.

two separate effects. First, the dynamics can distort local neighborhoods and probability-mass over time. Second, the geometry of the initial set determines whether every small spatial neighborhood contains enough volume, and hence enough probability-mass, to be detected from samples. This discussion motivates the following problem class. We require the flow to be well posed over the finite horizon of interest and Lipschitz with respect to initial conditions, so that local neighborhoods can be related through the pushforward dynamics. We also impose a geometric regularity condition on the initial set to exclude thin degeneracies. Throughout, for each finite horizon T considered below, we assume that $\varphi(t,\cdot)$ is well defined and unique on the relevant initial states for all $t \in [0,T]$; this is a standing well-posedness convention, separate from the statistical problem class. To state the geometric condition, we first recall the notion of reach.

Definition 1 (Reach [25])**.** *Let $S \subset \mathbb{R}^n$ be nonempty and closed. For $x \in \mathbb{R}^n$, define $d(x,S) := \inf_{y\in S} \|x-y\|$ and $\Pi_S(x) := \{y \in S : \|x-y\| = d(x,S)\}$. The reach of S is defined as follows:*

$$\operatorname{reach}(S) := \sup\{r \geq 0 : \Pi_S(x) \text{ is a singleton whenever } d(x,S) < r\}.$$

We now collect the dynamical and geometric requirements into a single problem family.

Definition 2 (Problem family)**.** *Fix constants $L > 0$ and $r_0 > 0$. Let $\mathcal{F}_{L,r_0}$ denote the class of triples (S_0, F, P_0) such that:*

1. *$S_0 \subset \mathcal{X}$ is a bounded open set with $0 < |S_0| < \infty$;*
2. *$\operatorname{reach}(S_0^c) \geq r_0$, where $S_0^c := \mathbb{R}^n \setminus S_0$;*
3. *$F : \mathcal{X} \to \mathbb{R}^n$ is L-Lipschitz on $\mathcal{X}$, i.e., $\|F(x) - F(y)\| \leq L\|x-y\|, \forall x, y \in \mathcal{X}$;*
4. *the initial distribution P_0 satisfies $\operatorname{supp}(P_0) = \overline{S_0}$.*

The problem family allows arbitrary finite positive volume. Throughout the paper, we use the normalization $|S_0| = 1$ only to fix the spatial scale and simplify constants. This normalization is not essential and can be replaced by any other fixed finite volume, with the corresponding rescaling of the reach parameter. This family gives a local thickness property: by [26, Lemma 4.8], $\overline{S_0} = \overline{S_0} \circ \mathcal{B}_{r_0}(0)$, positive reach of the complement rules out inward cusps and thin degeneracies, yielding lower bounds of the form $|\mathcal{B}_\rho(x) \cap \overline{S_0}| \gtrsim \rho^n$ for some small $\rho > 0$.[1] This local volume bound is the bridge from probability to geometry. After accounting for Lipschitz flow distortion and endpoint density lower bounds, it ensures that missing an r-scale region of the reachable set incurs non-negligible probability mass, which enables Hausdorff accuracy guarantees.

3 Reachable Set Approximation under Regular Geometry

Under the regularity conditions introduced above, we study Hausdorff approximation of the reachable set from finite endpoint samples. The key ingredient is a local endpoint probability lower bound, obtained by combining positive reach, an initial density lower bound, and Lipschitz flow distortion. This bound yields an inner-coverage sample-complexity result for any estimator containing all endpoint samples; a full Hausdorff guarantee follows after adding an estimator-dependent outer-deviation condition. We then prove that the r^{-n} dependence is information-theoretically unavoidable.

3.1 Local Mass Bounds under Flow Distortion

We first establish local mass lower bounds for the reachable set. Positive reach gives a static volume lower bound near the initial support, and the probability density lower bound converts it into sampling probability.

Proposition 1 (Lower Lebesgue mass bound for the initial set)**.** *Let $(S_0, F, P_0) \in \mathcal{F}_{L,r_0}$. Then, for every $x \in \overline{S_0}$ and every $r \in (0, r_0]$, we have $|\mathcal{B}_r(x) \cap \overline{S_0}| \geq \omega_n\, 2^{-n} r^n$.*

[1] $A \gtrsim B$ means $A \geq cB$ for a constant $c > 0$ independent of the scale of B.

Proof. See Appendix B.1. □

Assumption 1 (Initial density lower bound)**.** *The sampling distribution P_0 is absolutely continuous with respect to Lebesgue measure on $\overline{S_0}$, with density p_0. Moreover, there exists $\underline{p}_0 > 0$ such that $p_0(x) \geq \underline{p}_0$, for a.e. $x \in \overline{S_0}$.*

Combining Proposition 1 with Assumption 1 gives an initial local probability lower bound, making explicit the roles of r_0 and $\underline{p}_0$.

Corollary 1 (Initial local probability lower bound)**.** *Suppose $(S_0, F, P_0) \in \mathcal{F}_{L,r_0}$ and Assumption 1 holds. Then, for every $x \in \overline{S_0}$ and every $r \in (0, r_0]$,*

$$P_0(\mathcal{B}_r(x) \cap \overline{S_0}) \geq \underline{p}_0\, \omega_n\, 2^{-n} r^n.$$

To apply this local-mass argument to the reachable set at time T, we do not track the reach parameter itself. Instead, we propagate the local thickness implied by $\operatorname{reach}(S_0^c) \geq r_0$ through the Lipschitz flow and combine it with the propagated density lower bound.

Proposition 2 (Propagation of local Lebesgue mass)**.** *Let $(S_0, F, P_0) \in \mathcal{F}_{L,r_0}$. Then, for every $x \in \overline{S_T}$ and every $r \in (0, e^{LT} r_0]$, $|\mathcal{B}_r(x) \cap \overline{S_T}| \geq \omega_n\, 2^{-n} e^{-2nLT} r^n$.*

Proof. See Appendix B.3. □

Proposition 3 (Propagation of density lower bounds)**.** *Suppose $(S_0, F, P_0) \in \mathcal{F}_{L,r_0}$ and Assumption 1 holds. Then, P_T is absolutely continuous with respect to Lebesgue measure on $\overline{S_T}$. Moreover, if p_T denotes its density, then $p_T(y) \geq \underline{p}_0 e^{-nLT}$, for a.e. $y \in \overline{S_T}$. In particular, at the terminal time T, $\underline{p}_T := \underline{p}_0 e^{-nLT}$ is the sample density lower bound for the reachable set at time T.*

Proof. See Appendix B.3. □

Combining Proposition 2 with Proposition 3, we obtain the local endpoint probability lower bound used in the sampling-covering argument in the next subsection.

Corollary 2 (Endpoint local probability lower bound)**.** *Suppose $(S_0, F, P_0) \in \mathcal{F}_{L,r_0}$ and Assumption 1 holds. Define $\Gamma_T(r) := \underline{p}_0\, \omega_n 2^{-n} e^{-3nLT} r^n$. Then, for every $x \in \overline{S_T}$ and every $r \in (0, e^{LT} r_0]$,*

$$P_T(\mathcal{B}_r(x) \cap \overline{S_T}) \geq \Gamma_T(r).$$

Proof. See Appendix B.3. □

3.2 Sample Complexity Bounds

Corollary 2 shows that every r-neighborhood centered on the reachable set carries at least $\Gamma_T(r)$ endpoint probability mass. We now use this local probability lower bound to quantify how many endpoint samples are sufficient to cover the reachable set in the inner direction. The argument is estimator-agnostic: it only requires that the estimator contains all sampled endpoints.

Theorem 1 (Sampling upper bound for inner coverage)**.** *Let $(S_0, F, P_0) \in \mathcal{F}_{L,r_0}$, and $Y_1, \ldots, Y_N$ be* i.i.d. *samples from P_T. Suppose that Assumption 1 holds and the estimator $\widehat{S}_N$ contains all endpoint samples, i.e., $Y_i \in \widehat{S}_N$ for all $i = 1, \ldots, N$. Fix $0 < r \leq 2e^{LT} r_0$. If*

$$N \geq \frac{2^{2n} e^{3nLT}}{\underline{p}_0 \omega_n r^n} \left[\log \left(\frac{2^{3n} e^{nLT}}{\omega_n r^n} \right) + \log \frac{1}{\delta} \right],$$

then, with probability at least $1 - \delta$, $\sup_{x \in S_T} d(x, \widehat{S}_N) \leq r$.

Proof. See Appendix B.4. □

Corollary 3 (Hausdorff bound with outer deviation)**.** *Under the conditions of Theorem 1, suppose in addition that the estimator satisfies the outer-deviation condition* $\widehat{S}_N \subseteq \mathcal{B}_\eta(S_T)$ *for some* $\eta > 0$*. Then, with probability at least* $1-\delta$*,* $d_H(S_T, \widehat{S}_N) \le \max\{r, \eta\}$*.*

Proof. See Appendix B.4. □

Remark 1 (Estimator-dependent outer deviation)**.** *The parameter* η *measures the outer deviation of the estimator where* $\widehat{S}_N \subseteq \mathcal{B}_\eta(S_T)$ *is equivalent to* $\sup_{z\in\widehat{S}_N} d(z, S_T) \le \eta$*. In Corollary 3, the missed-mass condition controls the directed error from* S_T *to* $\widehat{S}_N$*, while* η *controls the opposite direction. Its value is estimator-dependent:* $\eta = 0$ *for inner estimators,* $\eta = h$ *for a union-of-balls estimator with radius* h*, and for a convex-hull estimator it is governed by the convexification error* $\sup_{z\in\mathrm{conv}(S_T)} d(z, S_T)$*. For a Christoffel-function estimator* $\widehat{S}_N^{\mathrm{Chr}} := \{z \in \mathcal{X} : \kappa_N(z) \le \tau_N\}$*, the corresponding value is* $\eta_{\mathrm{Chr,N}} := \sup_{\kappa_N(z)\le\tau_N} d(z, S_T)$*.*

We next show that the r^{-n} dependence is not a proof artifact, but an intrinsic limitation of sampling-based reachable-set approximation.

Theorem 2 (Minimax Hausdorff lower bound)**.** *Let* $C := 3^n\omega_n$*. For any accuracy* $0 < r \le 2^{-\frac{n+1}{n}} e^{LT} r_0$ *and any confidence level* $1-\delta$ *with* $\delta \in (0,1)$*, if* $N < \frac{e^{nLT}}{2^{n+1}Cr^n} \log\frac{1}{4\delta}$ *and* $r_0 \le C^{-\frac{1}{n}}$*, then for every estimator* $\widehat{S}_N$*, there exists an instance* $(S_0, F, P_0) \in \mathcal{F}_{L,r_0}$ *such that*

$$P_T^N(d_H(S_T, \widehat{S}_N) > r) > \delta.$$

Proof. See Appendix B.4. □

Together, the bounds show that the intrinsic dependence on the accuracy parameter is r^{-n}. The upper bound gives inner-coverage achievability up to logarithmic and flow-dependent factors, while the lower bound shows that $\Omega(e^{nLT}r^{-n}\log(1/\delta))$ samples are unavoidable in the worst case. Table 1 summarizes the comparison.

Table 1: Comparison between the upper and lower bounds.

	Upper bound	**Lower bound**
Type	Achievability	Minimax Hausdorff impossibility
Rate	$\mathcal{O}(e^{3nLT}r^{-n}[\log(1/\delta) + \log(e^{nLT}r^{-n})])$	$\Omega(e^{nLT}r^{-n}\log(1/\delta))$

4 Algorithm Design

The r^{-n} dependence in the lower bound highlights the intrinsic coverage burden of Hausdorff reachable-set approximation. This motivates the practical question studied in our experiments: under a fixed sample budget, can sampling be allocated more effectively than uniform sampling so as to improve geometric coverage and reduce the empirical burden of this curse of dimensionality? Accordingly, we focus on the sampling stage rather than on the downstream set representation. In the autonomous setting, samples are selected only through their initial conditions in S_0. For each sampled initial condition $X_j \in S_0$, we write $Y_j := \varphi(T, X_j)$ for the corresponding endpoint. Given endpoint samples $\mathcal{Y}_N = \{Y_j\}_{j=1}^N$, the reachable-set estimate is constructed as $\widehat{S}_N = \mathrm{C}(\mathcal{Y}_N)$, where $\mathrm{C}(\cdot)$ denotes a generic reachable-set estimator.

In the experiments, we compare uniform sampling with the adversarial sampling heuristic described by Lew and Pavone [19]. The purpose of using these two different sampling algorithms is to investigate whether adversarial sampling can improve finite-sample coverage by directing samples toward geometrically informative regions. We emphasize that the adversarial sampling procedure can improve finite sample coverage but cannot violate the lower bound. The worst-case Hausdorff sample complexity remains governed by the intrinsic r^{-n} coverage requirement. Implementation details of the adversarial sampling procedure (Algorithm 1) are provided in Appendix C.2. When $n_{\mathrm{adv}} = 0$ in Algorithm 1, no adversarial update is performed, and the method reduces to the uniform sampling method used in the experiments.

5 Experiments

In this section, we compare uniform and adversarial sampling on two 2D systems, one non-Lipschitz and one Lipschitz, and on a closed-loop robot-arm control task with different dimensions.

5.1 Adversarial Reachable Set Approximation under Non-Lipschitz Dynamics

Using the system in Figure 1, we first evaluate reachable-set approximation for the non-Lipschitz autonomous dynamics. This vector field is not globally Lipschitz and exhibits finite-time blow-up, so sampling errors can be strongly amplified in the positive x-direction. We consider three initial sets centered at $(2, 0)$: a disk, an equilateral triangle, and a morphologically opened triangle, all with the same area but different boundary geometry. For each set, we approximate the reachable set by the convex hull of propagated samples and report its Hausdorff distance to a high-resolution reference reachable set over time. Solid and dashed curves denote uniform and adversarial sampling, with shaded regions showing 95% confidence intervals over 50 random seeds.

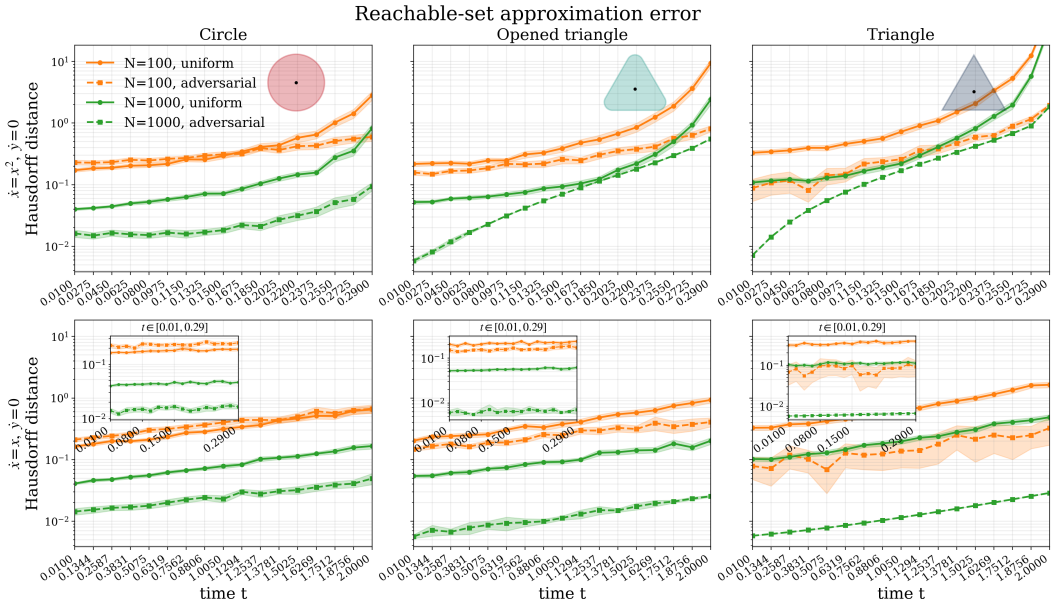

Figure 2: Hausdorff error versus time for $\dot{y} = 0, \dot{x} = x^2$ and $\dot{y} = 0, \dot{x} = x$ using different sampling methods. Solid and dashed curves denote uniform and adversarial sampling.

To isolate the effect of flow regularity, we also compare with the Lipschitz linear system $\dot{y} = 0, \dot{x} = x$ under the same initial-set configurations. Figure 2 shows that the linear system follows the expected exponential flow expansion, while the nonlinear system grows much faster due to finite-time blow-up. Across both systems, adversarial sampling reduces finite-sample error but mainly improves constants; the triangle and opened triangle remain harder than the disk, reflecting the effect of initial-set geometry.

5.2 Intrinsic Curse of Dimensionality Cannot be Circumvented

We also consider a vertical n-link robot arm tracking problem with $n \in \{2, 3, 4\}$. The state is $x = [q^\top, v^\top]^\top \in \mathbb{R}^{2n}$, where $q \in \mathbb{R}^n$ denotes joint angles and $v = \dot{q} \in \mathbb{R}^n$ denotes joint velocities. The trajectories are simulated in MuJoCo under the rigid-body dynamics $M(q)\dot{v} + C(q, v)v + g(q) = \tau$. Given a smooth reference trajectory $q_d(t)$, we use a non-adaptive tracking controller $\tau = M(q)\ddot{q}_d(t) + C(q, v)v + g(q) - K_p(q - q_d(t)) - K_d(v - \dot{q}_d(t))$ for control.

Although the nominal task is trajectory tracking, the initial state is uncertain, so safety verification requires characterizing how this uncertainty propagates through the closed-loop dynamics. We model the initial uncertainty by $S_0 = \left([-\rho_q, \rho_q]^n \times [-\rho_v, \rho_v]^n\right) \circ \mathcal{B}_r(0), \rho_q = \rho_v = 0.1, r = 0.01$.

We propagate sampled initial states to $T = 1.0$ second and measure approximation error by the point-cloud directed Hausdorff distance from a dense reference terminal cloud to the convex hull of the sampled terminal states.

As shown in Figure 3 and Table 2, the Hausdorff error decreases with the sample size under both uniform and adversarial sampling. However, the decay becomes slower as the state dimension increases: the fitted log–log slopes become less negative from dimension 4 (2-link) to dimension 8 (4-link). Adversarial sampling gives steeper slopes than uniform sampling in all dimensions, indicating better finite-sample efficiency. Nevertheless, its slopes also flatten as the dimension increases. Thus, targeted sampling improves the constants but does not remove the intrinsic dimension dependent degradation predicted by the r^{-n}-type sample-complexity scaling. Additional implementation details and more detailed comparisons and explanation, including the sample-budget-dependent improvement of adversarial sampling, are provided in Appendix C.3.

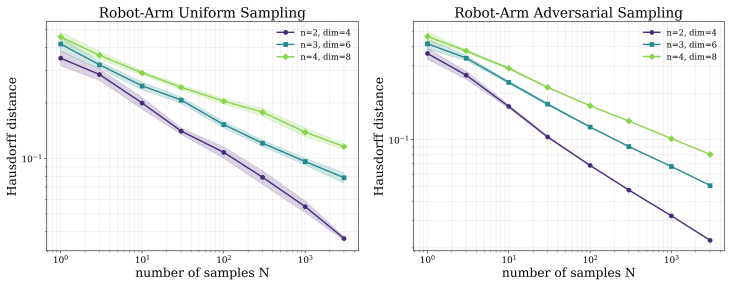

Figure 3: Hausdorff error versus sample size for robot-arm uncertainty propagation corresponding to different state dimensions under uniform sampling and adversarial sampling.

Dimension	4	6	8
Uniform sampling	-0.2806	-0.2094	-0.1662
Adversarial sampling	-0.3535	-0.2701	-0.2224

Table 2: Empirical log–log slopes for robot-arm uncertainty propagation. Each slope is obtained by fitting $\log d_H$ versus $\log N$, where d_H is the Hausdorff distance and N is the number of samples.

6 Conclusion and Future Work

This paper studied the finite-sample limits of sampling-based reachable-set approximation in Hausdorff distance. By viewing endpoint samples as observations from a pushforward distribution, we formulated reachable-set recovery as a geometric support estimation problem. We showed that, under positive-reach regularity of the initial domain and Lipschitz continuity of the autonomous flow, probability-mass coverage can be upgraded to Hausdorff accuracy. Our results give matching upper- and lower-bound perspectives. The upper bound shows that Hausdorff accuracy is attainable once every relevant r-scale region of the reachable set receives samples. The minimax lower bound shows that the dominant r^{-n} dependence, together with the flow-induced factor e^{nLT}, is intrinsic in the worst case. Experiments further support this picture: adversarial sampling improves finite-sample performance by emphasizing geometrically informative regions, but cannot remove the intrinsic dimension-dependent scaling. A natural direction for future work is to go beyond geometric regularity and global Lipschitz continuity by exploiting physics-informed inductive biases in the dynamics [27], such as energy conservation [28], contraction [29], or Hamiltonian structure [30], which may reduce the effective sample complexity for physically structured systems [31].

7 Limitations

The analysis is restricted to autonomous dynamics. This includes uncontrolled systems and closed-loop systems, but not open-loop reachable sets, which are unions over a control class rather than images of S_0 under a single flow map. As a result, the injectivity, volume-distortion, and endpoint-density arguments used here do not directly apply. Moreover, endpoint mass depends on how control signals are sampled, so geometrically important regions may receive very small probability mass.

References

[1] P. D. Holmes, S. Kousik, B. Zhang, D. Raz, C. Barbalata, M. Johnson-Roberson, and R. Vasudevan. Reachable sets for safe, real-time manipulator trajectory design. In *Proceedings of Robotics: Science and Systems*, 2020.

[2] Y. Meng, D. Sun, Z. Qiu, M. T. B. Waez, and C. Fan. Learning density distribution of reachable states for autonomous systems. In *Conference on Robot Learning*, pages 124–136. PMLR, 2022.

[3] R. Ivanov, T. Carpenter, J. Weimer, R. Alur, G. Pappas, and I. Lee. Verisig 2.0: Verification of neural network controllers using taylor model preconditioning. In *International Conference on Computer Aided Verification*, pages 249–262. Springer, 2021.

[4] J. Wu, H. Zhang, and Y. Vorobeychik. Verified safe reinforcement learning for neural network dynamic models. *Advances in Neural Information Processing Systems*, 37:117762–117783, 2024.

[5] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin. Hamilton-jacobi reachability: A brief overview and recent advances. In *2017 IEEE 56th annual conference on decision and control (CDC)*, pages 2242–2253. IEEE, 2017.

[6] M. Chen and C. J. Tomlin. Hamilton–jacobi reachability: Some recent theoretical advances and applications in unmanned airspace management. *Annual Review of Control, Robotics, and Autonomous Systems*, 1(1):333–358, 2018.

[7] A. Girard. Reachability of uncertain linear systems using zonotopes. In *Hybrid Systems: Computation and Control*, 2005.

[8] A. B. Kurzhanski and P. Varaiya. Ellipsoidal techniques for reachability analysis. In *Hybrid Systems: Computation and Control*, 2000.

[9] M. Selim, A. Alanwar, S. Kousik, G. Gao, M. Pavone, and K. H. Johansson. Safe reinforcement learning using black-box reachability analysis. *IEEE Robotics and Automation Letters*, 7(4): 10665–10672, 2022.

[10] M. Ganai, Z. Gong, C. Yu, S. Herbert, and S. Gao. Iterative reachability estimation for safe reinforcement learning. *Advances in Neural Information Processing Systems*, 36:69764–69797, 2023.

[11] J. Liu and E. Mallada. Recurrent control barrier functions: A path towards nonparametric safety verification. In *2025 IEEE 64th Conference on Decision and Control (CDC)*, pages 7721–7727. IEEE, 2025.

[12] Z. Ouyang, J. Liu, and E. Mallada. Symplectic inductive bias for data-driven target reachability in hamiltonian systems. *arXiv preprint arXiv:2604.17213*, 2026.

[13] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou. Information-theoretic model predictive control: Theory and applications to autonomous driving. *IEEE Transactions on Robotics*, 34(6):1603–1622, 2018.

[14] C. Pan, Z. Yi, G. Shi, and G. Qu. Sampling-based methods for optimal control: Theory, algorithms, and applications. In *2025 IEEE 64th Conference on Decision and Control (CDC)*, pages 3775–3793. IEEE, 2025.

[15] A. Castellano, S. Pan, and E. Mallada. Data-driven acceleration of mpc with guarantees. *arXiv preprint arXiv:2511.13588*, 2025.

[16] L. Liebenwein, C. Baykal, I. Gilitschenski, S. Karaman, and D. Rus. Sampling-based approximation algorithms for reachability analysis with provable guarantees. In *Proceedings of Robotics: Science and Systems*, 2018.

[17] A. Devonport, F. Yang, L. El Ghaoui, and M. Arcak. Data-driven reachability analysis with christoffel functions. In *2021 60th IEEE Conference on Decision and Control (CDC)*, pages 5067–5072. IEEE, 2021.

[18] A. Devonport, F. Yang, L. El Ghaoui, and M. Arcak. Data-driven reachability and support estimation with christoffel functions. *IEEE Transactions on Automatic Control*, 68(9):5216–5229, 2023.

[19] T. Lew and M. Pavone. Sampling-based reachability analysis: A random set theory approach with adversarial sampling. In *Conference on Robot Learning*, pages 2055–2070. PMLR, 2021.

[20] T. Lew, L. Janson, R. Bonalli, and M. Pavone. A simple and efficient sampling-based algorithm for general reachability analysis. In *Learning for Dynamics and Control Conference*, pages 1086–1099. PMLR, 2022.

[21] L. Devroye and G. L. Wise. Detection of abnormal behavior via nonparametric estimation of the support. *SIAM Journal on Applied Mathematics*, 38(3):480–488, 1980.

[22] A. Cuevas. Set estimation: Another bridge between statistics and geometry. *Boletín de Estadística e Investigación Operativa*, 25(2):71–85, 2009.

[23] A. P. Korostelev and A. B. Tsybakov. *Minimax Theory of Image Reconstruction*, volume 82 of *Lecture Notes in Statistics*. Springer, New York, 1993.

[24] E. Arias-Castro, B. Pateiro-López, and A. Rodríguez-Casal. Minimax estimation of the volume of a set under the rolling ball condition. *Journal of the American Statistical Association*, 114 (527):1162–1173, 2019.

[25] H. Federer. Curvature measures. *Transactions of the American Mathematical Society*, 93(3): 418–491, 1959.

[26] J. Rataj and M. Zähle. *Curvature measures of singular sets*. Springer, 2019.

[27] M. Raissi, P. Perdikaris, and G. E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378:686–707, 2019.

[28] T. Duong and N. Atanasov. Hamiltonian-based neural ode networks on the se(3) manifold for dynamics learning and control. In *Proceedings of Robotics: Science and Systems*, July 2021.

[29] S. Singh, B. Landry, A. Majumdar, J.-J. Slotine, and M. Pavone. Robust feedback motion planning via contraction theory. *The International Journal of Robotics Research*, 42(9):655–688, 2023.

[30] S. Greydanus, M. Dzamba, and J. Yosinski. Hamiltonian neural networks. *Advances in neural information processing systems*, 32, 2019.

[31] F. Djeumou, C. Neary, and U. Topcu. How to learn and generalize from three minutes of data: Physics-constrained and uncertainty-aware neural stochastic differential equations. In *Proceedings of The 7th Conference on Robot Learning*, pages 577–601. PMLR, 2023.

[32] G. C. Calafiore and M. C. Campi. The scenario approach to robust control design. *IEEE Transactions on Automatic Control*, 51(5):742–753, 2006.

[33] M. C. Campi and S. Garatti. The exact feasibility of randomized solutions of uncertain convex programs. *SIAM Journal on Optimization*, 19(3):1211–1230, 2008.

[34] A. Devonport and M. Arcak. Estimating reachable sets with scenario optimization. In *Proceedings of the 2nd Conference on Learning for Dynamics and Control*, volume 120 of *Proceedings of Machine Learning Research*, pages 75–84. PMLR, 2020.

[35] E. Dietrich, R. Devonport, S. Tu, and M. Arcak. Data-driven reachability with scenario optimization and the holdout method. In *2025 IEEE 64th Conference on Decision and Control (CDC)*, pages 3925–3931. IEEE, 2025.

[36] H. Sartipizadeh, A. P. Vinod, B. Açikmeşe, and M. Oishi. Voronoi partition-based scenario reduction for fast sampling-based stochastic reachability computation of linear systems. In *2019 American Control Conference (ACC)*, pages 37–44. IEEE, 2019.

[37] E. Dietrich, A. Devonport, and M. Arcak. Nonconvex scenario optimization for data-driven reachability. In *6th Annual Learning for Dynamics & Control Conference*, pages 514–527. PMLR, 2024.

[38] L. Hewing and M. N. Zeilinger. Scenario-based probabilistic reachable sets for recursively feasible stochastic model predictive control. *IEEE Control Systems Letters*, 4(2):450–455, 2019.

[39] S. Dean, N. Matni, B. Recht, and V. Ye. Robust guarantees for perception-based control. In *Proceedings of the 2nd Annual Conference on Learning for Dynamics and Control*, volume 120 of *Proceedings of Machine Learning Research*, pages 350–360. PMLR, 2020.

[40] C. Fan, B. Qi, S. Mitra, and M. Viswanathan. DryVR: Data-driven verification and compositional reasoning for automotive systems. In *Computer Aided Verification*, pages 441–461. Springer, 2017.

[41] S. A. Gruenbacher, M. Lechner, R. Hasani, D. Rus, T. A. Henzinger, S. A. Smolka, and R. Grosu. Gotube: Scalable statistical verification of continuous-depth models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 6755–6764, 2022.

[42] R. T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems*, volume 31, 2018.

[43] M. Fazlyab, A. Robey, H. Hassani, M. Morari, and G. J. Pappas. Efficient and accurate estimation of lipschitz constants for deep neural networks. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

[44] F. Latorre, P. Rolland, and V. Cevher. Lipschitz constant estimation of neural networks via sparse polynomial optimization. In *International Conference on Learning Representations*, 2020.

[45] G. Shafer and V. Vovk. A tutorial on conformal prediction. *Journal of Machine Learning Research*, 9:371–421, 2008.

[46] L. Lindemann, Y. Zhao, X. Yu, G. J. Pappas, and J. V. Deshmukh. Formal verification and control with conformal prediction: Practical safety guarantees for autonomous systems. *IEEE Control Systems*, 45(6):72–122, 2025.

[47] Z. Mei, A. Dixit, M. Booker, E. Zhou, M. Storey-Matsutani, A. Z. Ren, O. Shorinwa, and A. Majumdar. Perceive with confidence: Statistical safety assurances for navigation with learning-based perception. *The International Journal of Robotics Research*, 45(6):938–967, 2026.

[48] J. Yao, X. Zhang, Y. Xia, Z. Wang, A. Roy-Chowdhury, and J. Li. Towards generalizable safety in crowd navigation via conformal uncertainty handling. In *9th Annual Conference on Robot Learning*, 2025.

[49] I. Tabbara, Y. Yang, and H. Sibai. Statistically assuring safety of control systems using ensembles of safety filters and conformal prediction. *arXiv preprint arXiv:2511.07899*, 2025.

[50] L. Lindemann, M. Cleaveland, G. Shim, and G. J. Pappas. Safe planning in dynamic environments using conformal prediction. *IEEE Robotics and Automation Letters*, 8(8):5116–5123, 2023.

[51] A. Tebjou, G. Frehse, and F. Chamroukhi. Data-driven reachability using christoffel functions and conformal prediction. In *Proceedings of the Twelfth Symposium on Conformal and Probabilistic Prediction with Applications*, volume 204 of *Proceedings of Machine Learning Research*, pages 194–213. PMLR, 2023.

[52] S. Bansal and C. J. Tomlin. Deepreach: A deep learning approach to high-dimensional reachability. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1817–1824. IEEE, 2021.

[53] A. Lin and S. Bansal. Verification of neural reachable tubes via scenario optimization and conformal prediction. In *6th Annual Learning for Dynamics & Control Conference*, pages 719–731. PMLR, 2024.

[54] N. Hashemi, X. Qin, L. Lindemann, and J. V. Deshmukh. Data-driven reachability analysis of stochastic dynamical systems with conformal inference. In *2023 62nd IEEE Conference on Decision and Control (CDC)*, pages 3102–3109. IEEE, 2023.

[55] N. Hashemi, L. Lindemann, and J. V. Deshmukh. Pca-ddreach: Efficient statistical reachability analysis of stochastic dynamical systems via principal component analysis. In *Proceedings of the International Conference on Neuro-symbolic Systems*, volume 288, pages 693–707. PMLR, 2025.

[56] F. Bullo. *Contraction Theory for Dynamical Systems*. Kindle Direct Publishing, 1.2 edition, 2024. ISBN 979-8836646806. URL https://fbullo.github.io/ctds.

[57] L. C. Evans. *Measure theory and fine properties of functions*. Chapman and Hall/CRC, 2025.

[58] G. B. Folland. *Real analysis: modern techniques and their applications*. John Wiley & Sons, 1999.

[59] T. M. Cover. *Elements of information theory*. John Wiley & Sons, 1999.

[60] A. B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer, 1st edition, 2008.

[61] I. Goodfellow, J. Shlens, and C. Szegedy. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations*, 2015.

[62] Y. Dong, F. Liao, T. Pang, H. Su, J. Zhu, X. Hu, and J. Li. Boosting adversarial attacks with momentum. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9185–9193, 2018.

[63] A. Rodriguez-Casal and P. Saavedra-Nieves. A data-adaptive method for estimating density level sets under shape conditions. *The Annals of Statistics*, 50(3):1653–1668, 2022.

Appendix

The Appendix is organized as follows. Section A provides additional related work on sampling-based reachable set approximation methods. Section B presents the proofs of the theoretical results in the main paper, including the lower Lebesgue-mass bound, the mass-to-Hausdorff conversion argument, and sample complexity's lower bound and upper bound. Section C reports additional numerical experiments and implementation details.

Contents

A Further Related Work

Beyond the Hausdorff-oriented works discussed in the main paper, much of the sampling-based reachability literature studies weaker probabilistic notions of approximation, such as volume coverage, and violation probability, which are often characterized by scenario optimization and conformal prediction. We briefly review these adjacent frameworks below. While some of them provide important finite-sample and high-confidence guarantees, they do not by themselves ensure geometric recovery of the full reachable set. This makes them complementary to our goal of understanding when probability-level guarantees can be upgraded to Hausdorff accuracy.

A.1 Other Sampling-based Reachable Set Approximation Methods

Scenario optimization and sampling-based certificates. Scenario optimization provides a principled framework for deriving probabilistic guarantees from finitely many sampled constraints. In the classical scenario approach, an infinite family of constraints is replaced by a finite set of sampled constraints, leading to high-confidence bounds on the probability of constraint violation [32, 33]. This viewpoint has been adopted in reachability and safety verification, where sampled reachable states are used to construct data-driven outer approximations or safety certificates. For instance, scenario-based methods have been used to estimate reachable sets within tractable templates such as boxes, ellipsoids, and other parameterized set families [34, 35]. Related work has also studied reach-avoid problems for discrete-time linear time-invariant systems with additive probabilistic uncertainty, using Voronoi partition-based scenario reduction to accelerate sampling-based stochastic reachability computation [36]. Scenario-based techniques have also been extended to nonconvex set representations [37] and model-predictive control settings [38].These methods are attractive since they reduce reachable-set estimation or safety certification to finite-dimensional optimization problems with explicit confidence bounds. However, their guarantees are typically expressed in terms of violation probability or marginal coverage, rather than Hausdorff distance. This distinction is important: a set of small probability-mass can still be geometrically far from the estimator, especially in the presence of low-density regions, thin components, or separated reachable branches. Thus, such guarantees do not directly imply Hausdorff accuracy without additional geometric regularity and local lower-mass assumptions.

Set growth and Lipschitz constant-based bloating. Another line of sampling-based reachability methods estimates how sets grow under the dynamics by bounding the sensitivity of trajectories to initial conditions, disturbances, or model parameters. These methods often rely on Lipschitz constants, local discrepancy functions, finite-order Taylor approximations, or learned expansion rates to inflate sampled trajectories into tubes covering nearby trajectories. For example, some algorithms use smoothness of the dynamics together with Lipschitz constants to control the Lagrange remainder of Taylor approximations and thereby compute outer approximations of reachable sets [39]. Tools such as DryVR use discrepancy functions to quantify how the distance between trajectories evolves over time, allowing finitely many simulations to certify a larger reach tube [40]. More recent statistical verification methods, such as GoTube, estimate local expansion or Lipschitz-like quantities to construct probabilistic tubes for continuous-depth models and neural ODEs [41, 42]. A key limitation of this line of work is that tight Lipschitz constants are rarely available in data-driven or learning-based control systems. Computing upper bounds on the Lipschitz constants of neural networks is itself an active research problem [43, 44], and such bounds can be computationally expensive or overly conservative. Although scalable sampling-based procedures can estimate Lipschitz-like quantities, these estimates are generally not guaranteed to be valid upper bounds. These approaches are closely related to our flow-dependent analysis, since both use sensitivity of the dynamics to translate finite sampling resolution into geometric coverage after propagation. Our theory instead isolates the sample-complexity dependence needed for Hausdorff approximation of the reachable set itself, showing how flow distortion enters through dimension-dependent factors.

Conformal prediction and distribution-free coverage. Conformal prediction offers another route to finite-sample guarantees. Given a calibration dataset and a user-defined nonconformity score, conformal methods compute an empirical quantile of the calibration scores to construct prediction sets with distribution-free marginal coverage, where a future test sample is contained in the conformal set with a prescribed probability [45]. Since this construction does not require a parametric model of the data-generating distribution, it has been widely used as an uncertainty-quantification tool in learning-enabled autonomous systems, including perception, prediction, safe control, offline verification, and online monitoring [46, 47, 48, 49]. For instance, in safe planning for dynamic environments, conformal prediction calibrates prediction regions around learned trajectory predictors and incorporates them into model predictive control constraints to obtain user-specified probabilistic safety guarantees [50]. In reachability and safety verification, conformal prediction is typically used as a calibration layer on top of a data-driven reachable-set representation or a learned reachability surrogate. It has been used to calibrate Christoffel function's sublevel set thresholds for data-driven reachability, improving sample efficiency and robustness to outliers [51]; to verify DeepReach-style [52] neural HJ reachable tubes with probabilistic safety guarantees despite outlier errors [53]; and to inflate neural surrogate flowpipes for black-box stochastic systems using conformal residual bounds [54]. Recent PCA-based variants further reduce the conservatism of this inflation and improve scalability under distribution shift [55]. These methods are appealing because they require weak distributional assumptions and provide finite-sample coverage or safety guarantees. However, their guarantees remain probabilistic: a conformal set contains a future sample, trajectory, or prediction error with prescribed probability, but need not be close to the full reachable set in Hausdorff distance. It may cover most endpoint samples while missing a low-probability but spatially separated component. Thus, conformal prediction is complementary to our analysis, which identifies sampling conditions under which probability-mass coverage can be upgraded to Hausdorff accuracy.

B Proofs and Auxiliary Results

B.1 Lower Lebesgue Mass Bound for the Initial Set

We prove the local thickness property used in Proposition 1. Recall and that the positive-reach condition is imposed on the closed complement $S_0^c = \mathbb{R}^n \setminus S_0$.

Proposition 1 (Lower Lebesgue mass bound for the initial set). Let $(S_0, F, P_0) \in \mathcal{F}_{L,r_0}$. Then, for every $x \in \overline{S_0}$ and every $r \in (0, r_0]$,

$$|\mathcal{B}_r(x) \cap \overline{S_0}| \geq \omega_n \, 2^{-n} r^n.$$

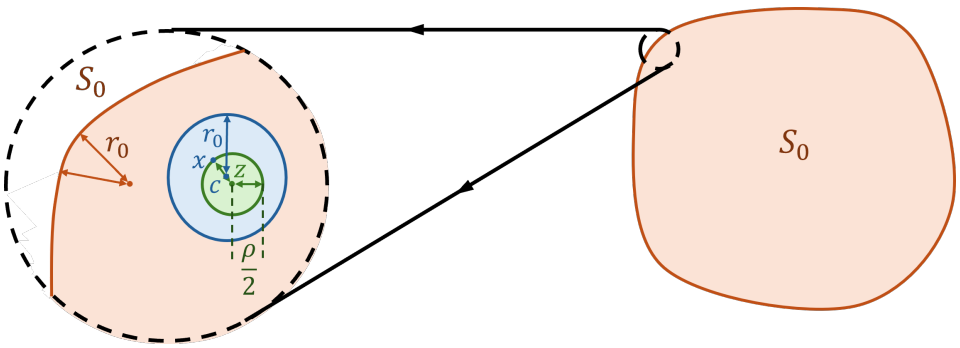

Figure 4: Local ball inclusion used in the proof of Proposition 1.

Proof of Proposition 1. Fix $x \in \overline{S_0}$ and $r \in (0, r_0]$. By the positive-reach opening property [26, Lemma 4.8], the condition $\operatorname{reach}(S_0^c) \geq r_0$ implies that $\overline{S_0}$ has an interior rolling-ball representation at scale r_0: for every $x \in \overline{S_0}$, there exists $c \in \overline{S_0}$ such that $x \in \mathcal{B}_{r_0}(c), \mathcal{B}_{r_0}(c) \subseteq \overline{S_0}$. We claim

that $\mathcal{B}_r(x) \cap \overline{S_0}$ contains a Euclidean ball of radius $r/2$. This immediately implies the desired lower bound. Let $d := \|x - c\|$. If $x = c$, then since $r \le r_0$,

$$\mathcal{B}_r(x) \subseteq \mathcal{B}_{r_0}(c) \subseteq \overline{S_0}.$$

Therefore,

$$|\mathcal{B}_r(x) \cap \overline{S_0}| = \omega_n r^n \ge \omega_n 2^{-n} r^n.$$

Now suppose $x \neq c$. Define $\nu := \frac{c-x}{\|c-x\|}$, $z := x + \frac{r}{2}\nu$. Then $\|z - x\| = r/2$, and hence $\mathcal{B}_{r/2}(z) \subseteq \mathcal{B}_r(x)$. It remains to show that $\mathcal{B}_{r/2}(z) \subseteq \mathcal{B}_{r_0}(c)$. Take any $w \in \mathcal{B}_{r/2}(z)$. If $d \ge r/2$, then $\|z - c\| = d - r/2$, and therefore

$$\|w - c\| \le \|w - z\| + \|z - c\| \le \frac{r}{2} + \left(d - \frac{r}{2}\right) = d \le r_0.$$

If $d < r/2$, then $\|z - c\| = r/2 - d$, and hence

$$\|w - c\| \le \|w - z\| + \|z - c\| \le \frac{r}{2} + \left(\frac{r}{2} - d\right) = r - d \le r \le r_0.$$

Thus, in both cases, $w \in \mathcal{B}_{r_0}(c)$. Since w was arbitrary,

$$\mathcal{B}_{r/2}(z) \subseteq \mathcal{B}_r(x) \cap \mathcal{B}_{r_0}(c) \subseteq \mathcal{B}_r(x) \cap \overline{S_0}.$$

Taking Lebesgue measure gives

$$|\mathcal{B}_r(x) \cap \overline{S_0}| \ge |\mathcal{B}_{r/2}(z)| = \omega_n \left(\frac{r}{2}\right)^n = \omega_n\, 2^{-n} r^n.$$

This proves the claim. □

B.2 Closure of the Reachable Set

We first record a simple topological fact used to identify the support of the endpoint distribution.

Lemma 1 (Reachability commutes with closure)**.** *Suppose $S_0 \subset \mathbb{R}^n$ is bounded and the finite-time flow map $\varphi(T, \cdot)$ is continuous on $\overline{S_0} = \overline{S_0}$. Then*

$$\mathcal{R}_T(\overline{S_0}) = \varphi(T, \overline{S_0}) = \overline{\varphi(T, S_0)} = \overline{\mathcal{R}_T(S_0)}.$$

In particular, if $S_T := \mathcal{R}_T(S_0)$, then $\mathcal{R}_T(\overline{S_0}) = \overline{S_T}$.

Proof. Since S_0 is dense in $\overline{S_0}$ and $\varphi(T, \cdot)$ is continuous on $\overline{S_0}$, we have $\varphi(T, \overline{S_0}) \subseteq \overline{\varphi(T, S_0)}$. Conversely, since S_0 is bounded, $\overline{S_0}$ is compact. Therefore $\varphi(T, \overline{S_0})$ is compact, hence closed. Moreover, $\varphi(T, S_0) \subseteq \varphi(T, \overline{S_0})$, so

$$\overline{\varphi(T, S_0)} \subseteq \varphi(T, \overline{S_0}).$$

Combining the two inclusions gives the claim. □

B.3 Volume Distortion and Density Propagation

We next prove the flow-regularity and density-propagation statements used in the main text. Throughout this subsection, $\overline{S_0} = \overline{S_0}$ and $\overline{S_T} := \mathcal{R}_T(\overline{S_0}) = \overline{\mathcal{R}_T(S_0)}$ by Lemma 1. We begin with the basic distortion properties of the finite-time flow. These estimates show that the flow is bi-Lipschitz on the relevant support and that Lebesgue volume cannot shrink by more than a factor e^{-nLT}.

Proposition 4 (Lipschitz regularity and geometry distortion of the flow)**.** *Let $(S_0, F, P_0) \in \mathcal{F}_{L,r_0}$. Then we have:*

1. *the map $x \mapsto \varphi(T, x)$ from $\overline{S_0}$ to $\overline{S_T}$ is injective;*

2. *for all*

$$x, y \in \overline{S_0}, \|\varphi(T,x) - \varphi(T,y)\| \le e^{LT}\|x - y\|;$$

3. *for all*

$$x, y \in \overline{S_0}, \|\varphi(T,x) - \varphi(T,y)\| \ge e^{-LT}\|x - y\|,$$

equivalently, the inverse map $\varphi(T,\cdot)^{-1} : \overline{S_T} \to \overline{S_0}$ *is* e^{LT}*-Lipschitz;*

4. *for every Lebesgue measurable set*

$$A \subseteq \overline{S_0}, |\mathcal{R}_T(A)| \ge e^{-nLT}|A|. \tag{3}$$

Proof of Proposition 4. We prove the four claims in such an order: 2 → 1 → 3 → 4. First, we prove the upper Lipschitz bound. Fix $x, y \in \overline{S_0}$. Since $(S_0, F, P_0) \in \mathcal{F}_{L,r_0}$, the vector field F is L-Lipschitz on $\mathcal{X}$. Hence, for all $s \in [0,T]$,

$$\|F(\varphi(s,x)) - F(\varphi(s,y))\| \le L\|\varphi(s,x) - \varphi(s,y)\|.$$

Using the integral form of the flow,

$$\varphi(s,x) - \varphi(s,y) = x - y + \int_0^s \left[F(\varphi(\sigma,x)) - F(\varphi(\sigma,y))\right] d\sigma.$$

Taking norms gives

$$\|\varphi(s,x) - \varphi(s,y)\| \le \|x - y\| + L\int_0^s \|\varphi(\sigma,x) - \varphi(\sigma,y)\|\, d\sigma.$$

By Gronwall's inequality [56, Corollary 3.17],

$$\|\varphi(s,x) - \varphi(s,y)\| \le e^{Ls}\|x - y\|, \qquad s \in [0,T].$$

Taking $s = T$ proves

$$\|\varphi(T,x) - \varphi(T,y)\| \le e^{LT}\|x - y\|.$$

Next, we prove injectivity. Suppose that $\varphi(T,a) = \varphi(T,b)$ for some $a, b \in \overline{S_0}$. Define the reverse-time trajectories

$$\eta_a(s) := \varphi(T - s, a), \qquad \eta_b(s) := \varphi(T - s, b), \qquad s \in [0,T].$$

Then $\eta_a(0) = \eta_b(0)$, and both trajectories solve the reverse-time equation

$$\dot{\eta}(s) = -F(\eta(s)).$$

Since F is L-Lipschitz, the reverse-time vector field $-F$ is also L-Lipschitz. By uniqueness of solutions, $\eta_a(s) = \eta_b(s)$ for all $s \in [0,T]$. Evaluating at $s = T$ gives $a = b$. Thus $x \mapsto \varphi(T,x)$ is injective on $\overline{S_0}$.

We now prove the lower Lipschitz bound. Fix $x, y \in \overline{S_0}$ and set

$$p := \varphi(T,x), \qquad q := \varphi(T,y).$$

Apply the upper Lipschitz estimate to the reverse-time dynamics. Since the reverse flow from p and q over time T returns to x and y, respectively, we obtain

$$\|x - y\| \le e^{LT}\|p - q\| = e^{LT}\|\varphi(T,x) - \varphi(T,y)\|.$$

Equivalently,

$$\|\varphi(T,x) - \varphi(T,y)\| \ge e^{-LT}\|x - y\|.$$

This also shows that the inverse map $\varphi(T,\cdot)^{-1} : \overline{S_T} \to \overline{S_0}$ is e^{LT}-Lipschitz.

It remains to prove the volume distortion bound. By [57, Theorem 2.8], we know that the inverse flow map $\varphi(T,\cdot)^{-1} : \overline{S_T} \to \overline{S_0}$, which is e^{LT}-Lipschitz, gives $|\varphi(T,\cdot)^{-1}(B)| \le e^{nLT}|B|, \forall B \subseteq \varphi(T, S_0)$ for every Lebesgue measurable set B. Taking $B = \varphi(T, A) = \mathcal{R}_t(A)$ and using injectivity of $x \mapsto \varphi(T, x)$ on $\overline{S_0}$, we have

$$\begin{aligned} |A| &= |\varphi(T,\cdot)^{-1}(\varphi(T, A))| \\ &\le e^{nLT}|\mathcal{R}_T(A)|. \end{aligned}$$

Rearranging yields $|\mathcal{R}_T(A)| \ge e^{-nLT}|A|$, which proves (3). This completes the proof. □

The preceding proposition is the technical tool needed to propagate the local thickness of the initial set. We now combine its Lipschitz and volume-distortion bounds with the initial local Lebesgue-mass bound from Proposition 1. This yields a corresponding local Lebesgue-mass lower bound for the reachable set at time T.

Proposition 2 (Propagation of local Lebesgue mass). Let $(S_0, F, P_0) \in \mathcal{F}_{L,r_0}$. Then, for every $x \in \overline{S_T}$ and every $r \in (0, e^{LT} r_0]$, $|\mathcal{B}_r(x) \cap \overline{S_T}| \ge \omega_n\, 2^{-n} e^{-2nLT} r^n$.

Proof of Proposition 2. Fix $x \in \overline{S_T}$ and $r \in (0, e^{LT} r_0]$. Since $\overline{S_T} = \varphi(T, \overline{S_0})$ and the flow map is injective on $\overline{S_0}$ by Proposition 4, there exists a unique $y \in \overline{S_0}$ such that $x = \varphi(T, y)$. Set $\rho := e^{-LT} r$. Then $0 < \rho \le r_0$. By Proposition 1,

$$|\mathcal{B}_\rho(y) \cap \overline{S_0}| \ge \omega_n 2^{-n} \rho^n.$$

Define

$$A := \mathcal{B}_\rho(y) \cap \overline{S_0}.$$

For every $p \in A$, the Lipschitz estimate in Proposition 4 gives

$$\|\varphi(T, p) - x\| = \|\varphi(T, p) - \varphi(T, y)\| \le e^{LT}\|p - y\| \le e^{LT}\rho = r.$$

Therefore, $\mathcal{R}_T(A) \subseteq \mathcal{B}_r(x) \cap \overline{S_T}$. Using the volume distortion lower bound in Proposition 4, we obtain $|\varphi(T, A)| \ge e^{-nLT}|A|$. Hence,

$$\begin{aligned} |\mathcal{B}_r(x) \cap \overline{S_T}| &\ge |\varphi(T, A)| \\ &\ge e^{-nLT}|A| \\ &\ge e^{-nLT}\omega_n 2^{-n}\rho^n \\ &= \omega_n 2^{-n} e^{-2nLT} r^n, \end{aligned}$$

where the last equality uses $\rho = e^{-LT} r$. □

The previous result propagates geometric thickness from $\overline{S_0}$ to $\overline{S_T}$. We also need to propagate the sampling density lower bound. The next proposition shows that the pushforward distribution remains absolutely continuous and that its density lower bound degrades by at most the same volume-distortion factor.

Proposition 3 (Propagation of density lower bounds). Suppose $(S_0, F, P_0) \in \mathcal{F}_{L,r_0}$ and Assumption 1 holds. Then $P_T = (\varphi(T,\cdot))_\# P_0$ is absolutely continuous with respect to Lebesgue measure on $\overline{S_T}$. Moreover, if p_T denotes its density, then

$$p_T(y) \ge \underline{p}_0 e^{-nLT}, \qquad \text{for a.e. } y \in \overline{S_T}.$$

In particular,

$$\underline{p}_T := \underline{p}_0 e^{-nLT}$$

is a density lower bound for the endpoint distribution at time T.

Proof of Proposition 3. Let $B \subseteq \overline{S_T}$ be a Lebesgue measurable set and define

$$A := \varphi(T,\cdot)^{-1}(B) \subseteq \overline{S_0}.$$

Since P_T is the pushforward of P_0 under $\varphi(T,\cdot)$,

$$P_T(B) = P_0(A).$$

We first prove absolute continuity. By Proposition 4, the inverse map $\varphi(T,\cdot)^{-1} : \overline{S_T} \to \overline{S_0}$ is e^{LT}-Lipschitz. Hence, for every Lebesgue measurable $B \subseteq \overline{S_T}$,

$$|A| = |\varphi(T,\cdot)^{-1}(B)| \le e^{nLT}|B|.$$

Therefore, if $|B| = 0$, then $|A| = 0$. Since P_0 is absolutely continuous with respect to Lebesgue measure on $\overline{S_0}$, this implies

$$P_T(B) = P_0(A) = 0.$$

Thus, P_T is absolutely continuous with respect to Lebesgue measure and admits a density p_T. It remains to prove the density lower bound. Since the flow map is injective on $\overline{S_0}$, we have $B = \mathcal{R}_T(A)$. By the upper Lipschitz bound in Proposition 4,

$$|B| = |\mathcal{R}_T(A)| \le e^{nLT}|A|.$$

Therefore, $|A| \ge e^{-nLT}|B|$. Using the lower bound on the initial density, we get

$$P_T(B) = P_0(A) = \int_A p_0(x)\,dx \ge \underline{p}_0|A| \ge \underline{p}_0 e^{-nLT}|B|.$$

Since this holds for every Lebesgue measurable set $B \subseteq \overline{S_T}$, the density p_T satisfies

$$p_T(y) \ge \underline{p}_0 e^{-nLT}, \qquad \text{for a.e. } y \in \overline{S_T}.$$

This proves the claim. □

We now combine the initial local Lebesgue-mass bound, the flow distortion estimate, and the endpoint density lower bound to prove the local probability lower bound used in the sample-complexity argument.

Corollary 2 (Endpoint local probability lower bound). Suppose $(S_0, F, P_0) \in \mathcal{F}_{L,r_0}$ and Assumption 1 holds. Define

$$\Gamma_T(r) := \underline{p}_0\,\omega_n 2^{-n} e^{-3nLT} r^n.$$

Then, for every $x \in \overline{S_T}$ and every $r \in (0, e^{LT}r_0]$,

$$P_T(\mathcal{B}_r(x) \cap \overline{S_T}) \ge \Gamma_T(r).$$

Proof of Corollary 2. By Proposition 3 and Proposition 3, we know for every $x \in \overline{S_T}$ and every $r \in (0, e^{LT}r_0]$, $|\mathcal{B}_r(x) \cap \overline{S_T}| \ge \omega_n\, 2^{-n} e^{-2nLT} r^n$ and the endpoint density satisfies $p_T(z) \ge \underline{p}_0 e^{-nLT}$, for a.e. $z \in \overline{S_T}$. Therefore,

$$\begin{aligned} P_T(\mathcal{B}_r(x) \cap \overline{S_T}) &= \int_{\mathcal{B}_r(x)\cap\overline{S_T}} p_T(z)\,dz \\ &\ge \underline{p}_0 e^{-nLT}|\mathcal{B}_r(x) \cap \overline{S_T}| \\ &\ge \underline{p}_0 e^{-nLT}\omega_n 2^{-n} e^{-2nLT} r^n \\ &= \underline{p}_0\,\omega_n 2^{-n} e^{-3nLT} r^n \\ &= \Gamma_T(r). \end{aligned}$$

This proves the claim. □

B.4 Sample Complexity

Sample Complexity Upper Bound. We prove Theorem 1 using a direct random-covering argument, which only uses the fact that $\widehat{S}_N$ contains all endpoint samples.

We first record the covering-number estimate used in the proof.

Definition 3 (Covering number)**.** *For a bounded set $A \subset \mathbb{R}^n$ and a radius $\rho > 0$, the ρ-covering number of A is defined as*

$$\mathcal{N}(A, \rho) := \min\{M \in \mathbb{N} : \exists z_1, \dots, z_M \in A \text{ such that } A \subseteq \bigcup_{j=1}^{M} \mathcal{B}_\rho(z_j)\}.$$

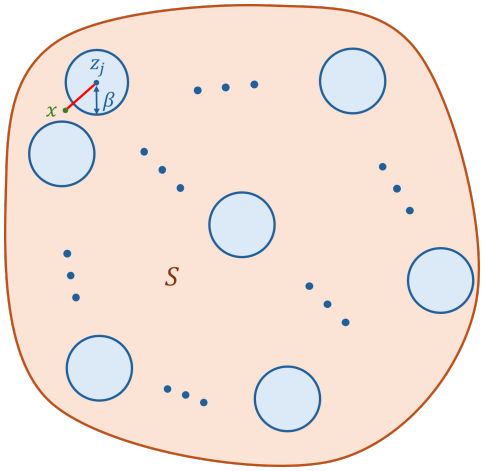

Figure 5: Set Covering

Lemma 2 (Covering number of the reachable set)**.** *Let $(S_0, F, P_0) \in \mathcal{F}_{L,r_0}$. For every $0 < \rho \leq e^{LT} r_0$,*

$$\mathcal{N}(\overline{S_T}, \rho) \leq \frac{2^{2n} e^{nLT}}{\omega_n \rho^n}.$$

In particular, for $\rho = r/2$, $\mathcal{N}(\overline{S_T}, r/2) \leq \frac{2^{3n} e^{nLT}}{\omega_n r^n}$.

Proof. Since the flow map $\varphi(T, \cdot)$ is e^{LT}-Lipschitz an injective on $\overline{S_0}$, any $e^{-LT}\rho$-cover of $\overline{S_0}$ is mapped to a ρ-cover of $\overline{S_T}$. Hence

$$\mathcal{N}(\overline{S_T}, \rho) \leq \mathcal{N}(\overline{S_0}, e^{-LT}\rho).$$

It remains to bound the covering number of $\overline{S_0}$. Fix $\beta \in (0, r_0]$. Let $\{z_1, \dots, z_M\} \subset \overline{S_0}$ be a maximal β-separated set, meaning that $\|z_i - z_j\| > \beta, i \neq j$, and no additional point of $\overline{S_0}$ can be added while preserving this property. By maximality, this set is a β-net of $\overline{S_0}$. Indeed, as Figure 5 shows, if there existed $x \in \overline{S_0}$ such that

$$x \notin \bigcup_{j=1}^{M} \mathcal{B}_\beta(z_j),$$

then $\|x - z_j\| > \beta$ for every j, so $\{z_1, \dots, z_M, x\}$ would still be β-separated, contradicting maximality. Therefore, $\overline{S_0} \subseteq \bigcup_{j=1}^{M} \mathcal{B}_\beta(z_j)$, and hence

$$\mathcal{N}(\overline{S_0}, \beta) \leq M.$$

We now upper bound M. Since the points $z_1, \dots, z_M$ are β-separated, the balls $\mathcal{B}_{\beta/2}(z_j)$ are pairwise disjoint up to their boundaries. Therefore, the sets

$$\mathcal{B}_{\beta/2}(z_j) \cap S_0, \qquad j = 1, \dots, M,$$

are also pairwise disjoint. By Proposition 1, applied with radius $\beta/2$, we have

$$\left|\mathcal{B}_{\beta/2}(z_j) \cap S_0\right| \geq \omega_n 2^{-n} \left(\frac{\beta}{2}\right)^n = \omega_n 2^{-2n} \beta^n, \qquad j = 1, \ldots, M.$$

Since these sets are pairwise disjoint subsets of S_0, and $|S_0| = 1$ in $\mathcal{F}_{L,r_0}$,

$$1 = |S_0| \geq \sum_{j=1}^{M} |\mathcal{B}_{\beta/2}(z_j) \cap S_0| \geq M\, \omega_n 2^{-2n} \beta^n.$$

Thus, $M \leq \frac{2^{2n}}{\omega_n \beta^n}$. Combining this with $\mathcal{N}(\overline{S_0}, \beta) \leq M$ gives

$$\mathcal{N}(\overline{S_0}, \beta) \leq \frac{2^{2n}}{\omega_n \beta^n}, \qquad 0 < \beta \leq r_0.$$

Finally, take $\beta = e^{-LT}\rho$. Since $\rho \leq e^{LT} r_0$, we have $\beta \leq r_0$, and therefore

$$\mathcal{N}(\overline{S_T}, \rho) \leq \mathcal{N}(\overline{S_0}, e^{-LT}\rho) \leq \frac{2^{2n}}{\omega_n (e^{-LT}\rho)^n} = \frac{2^{2n} e^{nLT}}{\omega_n \rho^n}.$$

Setting $\rho = r/2$ yields

$$\mathcal{N}(\overline{S_T}, r/2) \leq \frac{2^{3n} e^{nLT}}{\omega_n r^n}.$$

□

We now prove the sampling upper bound. The proof combines the endpoint local probability lower bound from Corollary 2 with a covering argument: we cover $\overline{S_T}$ by an $r/2$-net and show that, with high probability, every net ball is hit by at least one endpoint sample.

Theorem 1 (Sampling upper bound for inner coverage). Let $(S_0, F, P_0) \in \mathcal{F}_{L,r_0}$, and let $Y_1, \ldots, Y_N \overset{\text{i.i.d.}}{\sim} P_T$. Suppose that Assumption 1 holds and that the estimator $\widehat{S}_N$ contains all endpoint samples, i.e., $Y_i \in \widehat{S}_N$ for all $i = 1, \ldots, N$. Fix $0 < r \leq 2e^{LT} r_0$. If

$$N \geq \frac{2^{2n} e^{3nLT}}{\underline{p}_0 \omega_n r^n} \left[\log\left(\frac{2^{3n} e^{nLT}}{\omega_n r^n}\right) + \log\frac{1}{\delta}\right],$$

then, with probability at least $1 - \delta$,

$$\sup_{x \in S_T} d(x, \widehat{S}_N) \leq r.$$

Proof of Theorem 1. Let

$$q_r := \Gamma_T(r/2) = \underline{p}_0\, \omega_n 2^{-n} e^{-3nLT} \left(\frac{r}{2}\right)^n = \underline{p}_0\, \omega_n 2^{-2n} e^{-3nLT} r^n$$

and $\{z_1, \ldots, z_M\} \subset \overline{S_T}$ be an $r/2$-net of $\overline{S_T}$. By Lemma 2, we may choose this net so that $M \leq \frac{2^{3n} e^{nLT}}{\omega_n r^n}$. By Corollary 2, for each $j = 1, \ldots, M$,

$$P_T(\mathcal{B}_{r/2}(z_j) \cap \overline{S_T}) \geq \Gamma_T(r/2) = q_r.$$

Since P_T is supported on $\overline{S_T}$, this also gives $P_T(\mathcal{B}_{r/2}(z_j)) \geq q_r$. For each j, define the bad event

$$E_j := \left\{\mathcal{B}_{r/2}(z_j) \cap \{Y_1, \ldots, Y_N\} = \varnothing\right\}.$$

Let $\mathbb{P}(E_j)$ represents the probability that E_j happens. Because $Y_1, \ldots, Y_N$ are independent samples from P_T,

$$\mathbb{P}(E_j) = \left(1 - P_T(\mathcal{B}_{r/2}(z_j))\right)^N \leq (1 - q_r)^N \leq e^{-N q_r}.$$

Therefore, by the union bound,

$$\mathbb{P}\left(\bigcup_{j=1}^{M} E_j\right) \leq \sum_{j=1}^{M} \mathbb{P}(E_j) \leq M e^{-N q_r}.$$

Then suppose sample-size $N \geq \frac{2^{2n}e^{3nLT}}{\underline{p}_0\omega_n r^n}\left[\log\left(\frac{2^{3n}e^{nLT}}{\omega_n r^n}\right) + \log\frac{1}{\delta}\right]$, because $\frac{1}{q_r} = \frac{2^{2n}e^{3nLT}}{\underline{p}_0\omega_n r^n}$ and $M \leq \frac{2^{3n}e^{nLT}}{\omega_n r^n}$. Thus,$Me^{-Nq_r} \leq \delta$. Hence, with probability at least $1-\delta$, none of the events E_j occurs. Equivalently, every $r/2$-net ball contains at least one endpoint sample. On this event, fix any $x \in \overline{S_T}$. Since $\{z_1, \ldots, z_M\}$ is an $r/2$-net, there exists z_j such that $\|x - z_j\| \leq r/2$. Since E_j does not occur, there exists a sample Y_i such that $\|Y_i - z_j\| \leq r/2$. By the sample-containment assumption, $Y_i \in \widehat{S}_N$. Therefore,

$$d(x, \widehat{S}_N) \leq \|x - Y_i\| \leq \|x - z_j\| + \|z_j - Y_i\| \leq r.$$

Since this holds for every $x \in \overline{S_T}$, we conclude that

$$\sup_{x \in S_T} d(x, \widehat{S}_N) = \sup_{x \in \overline{S_T}} d(x, \widehat{S}_N) \leq r.$$

This completes the proof. □

The preceding theorem controls only the inner directed error from the reachable set to the estimator. We now add the estimator-dependent outer deviation to obtain a full Hausdorff bound.

Corollary 3 (Hausdorff bound with outer deviation). Under the conditions of Theorem 1, suppose in addition that the estimator satisfies the outer-deviation condition

$$\widehat{S}_N \subseteq \mathcal{B}_\eta(S_T)$$

for some $\eta > 0$. Then, with probability at least $1-\delta$,

$$d_H(S_T, \widehat{S}_N) \leq \max\{r, \eta\}.$$

Proof of Corollary 3. By Theorem 1, with probability at least $1-\delta$, $\sup_{x \in \overline{S_T}} d(x, \widehat{S}_N) \leq r$. The outer-deviation condition $\widehat{S}_N \subseteq \mathcal{B}_\eta(S_T)$ implies $\sup_{z \in \widehat{S}_N} d(z, S_T) \leq \eta$. Therefore, on the same event,

$$\begin{aligned} d_H(S_T, \widehat{S}_N) &= \max\{\sup_{x \in S_T} d(x, \widehat{S}_N), \sup_{z \in \widehat{S}_N} d(z, S_T)\} \\ &\leq \max\{r, \eta\}. \end{aligned}$$

Thus, with probability at least $1-\delta$,

$$d_H(S_T, \widehat{S}_N) \leq \max\{r, \eta\}.$$

□

Sample Complexity Lower Bound. To obtain the sample-complexity lower bound, we first identify the local statistical obstruction that any sampling-based method must face. The upper bound relies on the fact that every r-scale neighborhood of the reachable set has nontrivial probability-mass. For the lower bound, we need the converse type of phenomenon: among many disjoint r-scale regions inside a regular support, at least one region must have small probability-mass. Otherwise, the total probability over all such regions would exceed one.

We formalize this idea through an interior packing number.

Definition 4 (Interior packing number)**.** *Let $S_0 \subset \mathbb{R}^n$ be a nonempty set and let $r > 0$. The interior r-packing number of S_0, denoted by $M_r(S_0)$, is the largest integer M for which there exist points $z_1, \ldots, z_M \in S_0$ such that*

$$\mathcal{B}_r(z_j) \subset S_0, \qquad \mathcal{B}_r(z_i) \cap \mathcal{B}_r(z_j) = \emptyset \quad (i \neq j).$$

Equivalently, $M_r(S_0)$ is the maximum number of pairwise disjoint closed balls of radius r that can be placed entirely inside S_0.

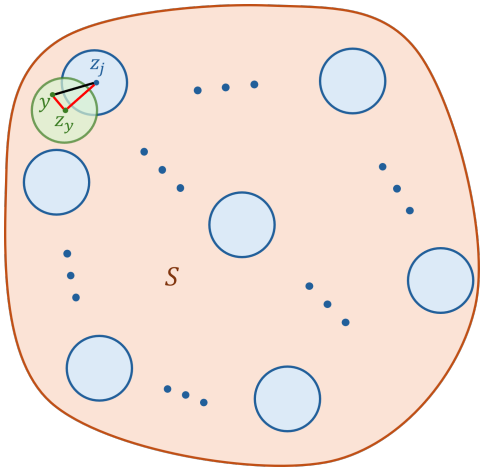

Figure 6: Interior Packing

Positive reach provides the geometric ingredient needed to lower bound this packing number. It ensures that the support has enough interior thickness to contain order r^{-n} disjoint balls of radius r. A simple averaging argument then shows that, for any probability measure supported on S_0, at least one of these balls has probability-mass at most order r^n.

Lemma 3 (Interior packing from positive reach)**.** *Let $S_0 \subset \mathbb{R}^n$ be a bounded open set satisfying* $\operatorname{reach}(S_0^c) \geq r_0 > 0$ *and* $|S_0| = 1$*. Then, for every $0 < r \leq r_0$, there exist points $z_1, \ldots, z_{M_r} \in S_0$ such that $\mathcal{B}_r(z_j) \subset S_0, \mathcal{B}_r(z_i) \cap \mathcal{B}_r(z_j) = \varnothing\ (i \neq j)$, and $M_r \geq \frac{3^{-n} r^{-n}}{\omega_n}$.*

Proof. Since $\operatorname{reach}(S_0^c) \geq r_0$, the opening property implies that $S_0 = S_0 \circ \mathcal{B}_r(0)$ for every $0 < r \leq r_0$. Hence, as Figure 6 shows, for any $r \in (0, r_0)$ and for every $y \in S_0$, there exists a point $z_y \in S_0 \ominus \mathcal{B}_r(0)$ such that $y \in \mathcal{B}_r(z_y)$. In particular,

$$\mathcal{B}_r(z_y) \subset S_0.$$

Let $\{\mathcal{B}_r(z_j)\}_{j=1}^{M_r}$ be a maximal family of pairwise disjoint balls contained in S_0. By maximality, for every $y \in S_0$, the interior ball $\mathcal{B}_r(z_y) \subset S_0$ must intersect at least one selected ball $\mathcal{B}_r(z_j)$; otherwise it could be added to the family. Thus $\|z_y - z_j\| \leq 2r$ for some j. Since $y \in \mathcal{B}_r(z_y)$, we get

$$\|y - z_j\| \leq \|y - z_y\| + \|z_y - z_j\| \leq 3r.$$

Therefore,

$$S_0 \subseteq \bigcup_{j=1}^{M_r} \mathcal{B}_{3r}(z_j).$$

Taking Lebesgue measure gives

$$1 = |S_0| \leq \sum_{j=1}^{M_r} |\mathcal{B}_{3r}(z_j)| = M_r \omega_n (3r)^n.$$

Rearranging yields

$$M_r \geq \frac{3^{-n} r^{-n}}{\omega_n}.$$

This proves the claim. □

The lemma above is purely geometric: it lower bounds the number of disjoint r-scale balls that can be placed inside a positive-reach set. We now turn this packing statement into a probabilistic one. Since the balls are disjoint and the total probability-mass is one, at least one of these balls must carry probability no larger than the average mass $1/M_r$.

Corollary 4 (Existence of a low-probability ball)**.** *Under the conditions of Lemma 3, for any probability measure P_0 supported on $\overline{S_0}$, there exists $j \in \{1, \dots, M_r\}$ such that*

$$P_0(\mathcal{B}_r(z_j)) \le \frac{1}{M_r} \le 3^n \omega_n r^n.$$

Proof. Since the balls $\mathcal{B}_r(z_1), \dots, \mathcal{B}_r(z_{M_r})$ are pairwise disjoint, $\bigcup_{j=1}^{M_r} \mathcal{B}_r(z_j) \subseteq \overline{S_0}$ and P_0 is a probability measure on the Borel σ-algebra of $\mathbb{R}^n$ supported on $\overline{S_0}$,

$$1 = P_0(\overline{S_0}) \ge P_0\left(\bigcup_{j=1}^{M_r} \mathcal{B}_r(z_j)\right) = \sum_{j=1}^{M_r} P_0(\mathcal{B}_r(z_j)).$$

Hence at least one term satisfies $P_0(\mathcal{B}_r(z_j)) \le 1/M_r$. Using Lemma 3 gives the fact that there exists $j \in \{1, \dots, M_r\}$ such that

$$P_0(\mathcal{B}_r(z_j)) \le \frac{1}{M_r} \le 3^n \omega_n r^n.$$

□

This low-probability ball will be removed in the two-point construction below, producing a Hausdorff-visible perturbation that is statistically difficult to detect. More precisely, we will construct two admissible reachability instances whose endpoint supports are close in distribution but far in Hausdorff distance. The two instances differ only through the initial domain and sampling law: the first uses the original initial domain, while the second removes the low-probability interior ball from it. After propagation through the flow, these two initial domains induce two endpoint supports.

To make the testing reduction independent of this particular construction, we first state it for two generic candidate endpoint supports.

Definition 5 (Hypothesis test induced by a support estimator)**.** *Let $S^{(0)}, S^{(1)} \subset \mathbb{R}^n$ be two candidate endpoint supports. Given a support estimator $\widehat{S}_N$, define the induced binary test $\psi = \psi(Y_1, \dots, Y_N) \in \{0, 1\}$ by*

$$\psi(Y_1, \dots, Y_N) := \begin{cases} 0, & d_H(\widehat{S}_N, S^{(0)}) \le d_H(\widehat{S}_N, S^{(1)}), \\ 1, & d_H(\widehat{S}_N, S^{(1)}) < d_H(\widehat{S}_N, S^{(0)}). \end{cases}$$

The test simply chooses the candidate endpoint support that is closer to the estimated set in Hausdorff distance. Therefore, if the estimator is accurate and the two candidate endpoint supports are separated by more than twice the target accuracy, the induced test must also be correct. This observation is formalized next.

Lemma 4 (Accurate support estimation induces a correct test)**.** *Suppose $d_H(S^{(0)}, S^{(1)}) > 2r$, and let ψ be the test in Definition 5. If the true endpoint support is $S^{(0)}$ and*

$$d_H(\widehat{S}_N, S^{(0)}) \le r,$$

then $\psi = 0$. Similarly, if the true endpoint support is $S^{(1)}$ and

$$d_H(\widehat{S}_N, S^{(1)}) \le r,$$

then $\psi = 1$.

Proof. If $d_H(\widehat{S}_N, S^{(0)}) \le r$, then by the triangle inequality,

$$d_H(\widehat{S}_N, S^{(1)}) \ge d_H(S^{(0)}, S^{(1)}) - d_H(\widehat{S}_N, S^{(0)}) > 2r - r = r.$$

Hence $d_H(\widehat{S}_N, S^{(0)}) < d_H(\widehat{S}_N, S^{(1)})$, and therefore $\psi = 0$ by Definition 5. The argument for $S^{(1)}$ is identical. □

To compare two such reachable-set instances statistically, we need to quantify how distinguishable the corresponding endpoint-sampling distributions are. This motivates the following two standard information-theoretic notions. The Radon–Nikodym derivative allows us to express one probability measure as a density with respect to another, and the Kullback–Leibler divergence then measures the information gap between the two induced sampling models.

Definition 6 (Radon–Nikodym derivative [58])**.** *Let P and Q be probability measures on a measurable space $(\Omega, \mathcal{F})$. We say that P is absolutely continuous with respect to Q, denoted by $P \ll Q$, if $Q(A) = 0$ implies $P(A) = 0$ for every $A \in \mathcal{F}$. In this case, the Radon–Nikodym derivative $\frac{dP}{dQ}$ is the measurable function satisfying*

$$P(A) = \int_A \frac{dP}{dQ}\, dQ, \qquad \forall A \in \mathcal{F}.$$

Definition 7 (Kullback–Leibler divergence [59])**.** *Let P and Q be probability measures on $(\Omega, \mathcal{F})$. If $P \ll Q$, the Kullback–Leibler divergence from P to Q is*

$$D_{\mathrm{KL}}(P\|Q) := \int_\Omega \log\left(\frac{dP}{dQ}\right) dP.$$

If $P \not\ll Q$, we set $D_{\mathrm{KL}}(P\|Q) = +\infty$.

The preceding lemma reduces Hausdorff estimation to binary testing. We now use a standard testing lower bound to show that this binary decision problem cannot be solved reliably when the two sampling distributions have small KL divergence.

Lemma 5 ([60], Lemma 2.6)**.** *Let P and Q be probability measures with $P \ll Q$. Then, for any binary test $\psi \in \{0, 1\}$,*

$$P(\psi = 1) + Q(\psi = 0) \geq \frac{1}{2} \exp(-D_{\mathrm{KL}}(P\|Q))\,.$$

Consequently, for N i.i.d. *samples,*

$$P^N(\psi = 1) + Q^N(\psi = 0) \geq \frac{1}{2} \exp\big(-D_{\mathrm{KL}}(P^N\|Q^N)\big)\,.$$

With these ingredients in place, we now prove the minimax lower bound. The proof constructs two admissible instances

$$(S_0^{(0)}, F^{(0)}, P_0^{(0)}) \quad \text{and} \quad (S_0^{(1)}, F^{(1)}, P_0^{(1)})$$

in $\mathcal{F}_{L,r_0}$. The second instance is obtained by removing a low-probability interior ball from the initial domain of the first. Under an admissible expanding flow, the resulting endpoint supports $\overline{S_T}^{(0)}$ and $\overline{S_T}^{(1)}$ are separated in Hausdorff distance, while the endpoint sampling distributions $P_T^{(0)}$ and $P_T^{(1)}$ have small KL divergence.

We now prove the minimax lower bound. The proof uses a two-point testing argument. We construct two admissible instances in the problem family whose endpoint supports are separated in Hausdorff distance, but whose endpoint sampling distributions are close in KL divergence. Any estimator that is Hausdorff-accurate on both instances would therefore induce a reliable binary test, contradicting the testing lower bound when the sample size is too small.

Theorem 2 (Minimax Hausdorff lower bound). Let $C := 3^n \omega_n$. For any accuracy $0 < r \leq 2^{-\frac{n+1}{n}} e^{LT} r_0$ and any confidence level $1 - \delta$ with $\delta \in (0, 1)$, if

$$N < \frac{e^{nLT}}{2^{n+1} C r^n} \log \frac{1}{4\delta}, \text{ and } r_0 \leq C^{-\frac{1}{n}}$$

then for every estimator $\widehat{S}_N$, there exists an instance $(S_0, F, P_0) \in \mathcal{F}_{L,r_0}$ such that

$$P_T^N\left(d_H(S_T, \widehat{S}_N) > r\right) > \delta.$$

Proof of Theorem 2. We prove the result by a two-point minimax testing argument. Fix an arbitrary estimator $\widehat{S}_N$. It suffices to construct two admissible L-Lipschitz dynamics and initial distributions whose endpoint supports are separated in Hausdorff distance, but whose endpoint sample distributions are statistically close. To be more specific, our goal is to construct two admissible instances $(S_0^{(0)}, F, P_0^{(0)})$ and $(S_0^{(1)}, F, P_0^{(1)})$ are in $\mathcal{F}_{L,r_0}$ whose endpoint supports, denoted by $\overline{S_T}^{(0)}$ and $\overline{S_T}^{(1)}$, are separated in Hausdorff distance, while their endpoint sampling distributions $P_T^{(0)}$ and $P_T^{(1)}$ are statistically close.

Let $R := \omega_n^{-1/n}$ and choose $x_0 \in \mathbb{R}^n$. Define the reference initial domain by $S_0^{(0)} := \mathcal{B}_R^\circ(x_0), \overline{S_0}^{(0)} := \overline{S_0^{(0)}} = \mathcal{B}_R(x_0), P_0^{(0)} := \mathrm{Unif}(S_0^{(0)})$. Then $|S_0^{(0)}| = 1$ and $\mathrm{supp}(P_0^{(0)}) = \overline{S_0}^{(0)}$. Moreover, since $R = \omega_n^{-1/n}$ and $r_0 \le C^{-1/n} = \frac{1}{3}\omega_n^{-1/n}$, we have $\mathrm{reach}((S_0^{(0)})^c) \ge r_0$ and $S_0^{(0)} \ominus \mathcal{B}_{3r_0}(0) \neq \varnothing$.

We use the admissible L-Lipschitz vector field $F(x) = Lx$, whose flow is $\varphi(T, x) = e^{LT}x$. In particular, $\varphi(T, \mathcal{B}_\rho(z)) = \mathcal{B}_{e^{LT}\rho}(\varphi(T, z)), \forall \rho > 0,\ z \in \mathbb{R}^n$. Fix $\alpha > 0$ sufficiently small and set $\rho := (2+\alpha)e^{-LT}r$. By the assumed upper range on r, we may choose α sufficiently small so that $\rho = (2+\alpha)e^{-LT}r \le r_0$. Applying Corollary 4 to $S_0^{(0)}$ at radius ρ, we obtain a point $z_j \in S_0^{(0)} \ominus \mathcal{B}_{3r_0}(0)$ such that $\mathcal{B}_\rho(z_j) \subset S_0^{(0)}$ and

$$q_r := P_0^{(0)}(\mathcal{B}_\rho^\circ(z_j)) \le P_0^{(0)}(\mathcal{B}_\rho(z_j)) \le C\rho^n = (2+\alpha)^n C e^{-nLT} r^n.$$

The condition $r \le 2^{-(n+1)/n} e^{LT} r_0$ together with $r_0 \le C^{-1/n}$ implies $q_r \le \frac{1}{2}(\frac{2+\alpha}{2})^n$.

The perturbed initial domain is defined by removing this closed ball: $S_0^{(1)} := S_0^{(0)} \setminus \mathcal{B}_\rho(z_j)$. Since $S_0^{(0)}$ is open and $\mathcal{B}_\rho(z_j)$ is closed, $S_0^{(1)}$ is open. Its corresponding support is $\overline{S_0}^{(1)} := \overline{S_0^{(1)}} = \overline{S_0}^{(0)} \setminus \mathcal{B}_\rho^\circ(z_j)$. We define the second initial distribution by conditioning on this new support: $P_0^{(1)} := P_0^{(0)}(\cdot \mid \overline{S_0}^{(1)})$. Equivalently, up to boundary sets of Lebesgue measure zero, $P_0^{(1)} = \mathrm{Unif}(S_0^{(1)})$, and hence $\mathrm{supp}(P_0^{(1)}) = \overline{S_0}^{(1)}$. It remains to verify the reach condition for the perturbed domain. Since

$$(S_0^{(1)})^c = (S_0^{(0)})^c \cup \mathcal{B}_\rho(z_j),$$

and $z_j \in S_0^{(0)} \ominus \mathcal{B}_{3r_0}(0)$ while $\rho \le r_0$, the two closed components $(S_0^{(0)})^c$ and $\mathcal{B}_\rho(z_j)$ are separated by distance at least $2r_0$. Moreover, $(S_0^{(0)})^c$ has reach at least r_0, and $\mathcal{B}_\rho(z_j)$ is convex. Therefore, by the separated-union reach lemma,

$$\mathrm{reach}((S_0^{(1)})^c) \ge r_0.$$

Thus both $(S_0^{(0)}, F, P_0^{(0)})$ and $(S_0^{(1)}, F, P_0^{(1)})$ belong to $\mathcal{F}_{L,r_0}$.

By the L-exponential growth property, $\varphi(T, \mathcal{B}_{2e^{-LT}r}(z_j)) = \mathcal{B}_{2r}(y_j)$. Thus a low-probability initial ball at scale $2e^{-LT}r$ becomes an endpoint ball at Hausdorff scale $2r$. We first verify the Hausdorff separation between the two endpoint supports. Recall that $\overline{S_T}^{(0)} := \varphi(T, \overline{S_0}^{(0)})$, $\overline{S_T}^{(1)} := \varphi(T, \overline{S_0}^{(1)})$, where $\overline{S_0}^{(1)} = \overline{S_0}^{(0)} \setminus \mathcal{B}_\rho^\circ(z_j)$. Since the flow is $\varphi(T, x) = e^{LT}x$ and $\rho = (2+\alpha)e^{-LT}r$, we have $\overline{S_T}^{(1)} = \overline{S_T}^{(0)} \setminus \mathcal{B}_{2r}^\circ(y_j)$, $y_j := \varphi(T, z_j)$. In particular, $y_j \in \overline{S_T}^{(0)}$ and the closest points of $\overline{S_T}^{(1)}$ to y_j lie on the boundary of $\mathcal{B}_{2r}(y_j)$. Hence

$$d(y_j, \overline{S_T}^{(1)}) = (2+\alpha)r > 2r. \tag{4}$$

Therefore, $d_H(\overline{S_T}^{(0)}, \overline{S_T}^{(1)}) \ge d(y_j, \overline{S_T}^{(1)}) > 2r$. Next, we bound the KL divergence between the two endpoint distributions. Since

$$P_0^{(1)} = P_0^{(0)}(\cdot \mid \overline{S_0}^{(1)}), \qquad P_0^{(0)}(\overline{S_0}^{(1)}) = 1 - q_r,$$

we have, for every measurable set $A \subseteq \mathbb{R}^n$,

$$P_0^{(1)}(A) = \frac{P_0^{(0)}(A \cap \overline{S_0}^{(1)})}{1 - q_r}.$$

Equivalently, this conditional distribution can be written in Radon–Nikodym form as $P_0^{(1)}(A) = \int_A \frac{\mathbf{1}_{\overline{S_0}^{(1)}}(x)}{1-q_r}\, dP_0^{(0)}(x)$. Therefore,

$$\frac{dP_0^{(1)}}{dP_0^{(0)}}(x) = \frac{\mathbf{1}_{\overline{S_0}^{(1)}}(x)}{1-q_r}, \qquad P_0^{(0)}\text{-a.s.}$$

We now compute the KL divergence. By definition,

$$D_{\mathrm{KL}}(P_0^{(1)} \| P_0^{(0)}) = \int \log\left(\frac{dP_0^{(1)}}{dP_0^{(0)}}(x)\right) dP_0^{(1)}(x).$$

Since $P_0^{(1)}$ is supported on $\overline{S_0}^{(1)}$, we have $\mathbf{1}_{\overline{S_0}^{(1)}}(x) = 1$ for $P_0^{(1)}$-almost every x. Hence, on the support of $P_0^{(1)}$,

$$\frac{dP_0^{(1)}}{dP_0^{(0)}}(x) = \frac{1}{1-q_r}.$$

Therefore,

$$\begin{aligned} D_{\mathrm{KL}}(P_0^{(1)} \| P_0^{(0)}) &= \int_{\overline{S_0}^{(1)}} \log\left(\frac{1}{1-q_r}\right) dP_0^{(1)}(x) \\ &= \log\left(\frac{1}{1-q_r}\right) \int_{\overline{S_0}^{(1)}} dP_0^{(1)}(x) \\ &= \log\left(\frac{1}{1-q_r}\right) P_0^{(1)}(\overline{S_0}^{(1)}) \\ &= \log\left(\frac{1}{1-q_r}\right) \\ &= -\log(1-q_r). \end{aligned} \tag{5}$$

Because $\varphi(T,\cdot)$ is a one-to-one flow map on the supports under consideration in the hard instance (e.g. $\dot{x} = Lx$), KL divergence is invariant under pushforward by $\varphi(T,\cdot)$. Thus

$$\begin{aligned} D_{\mathrm{KL}}(P_T^{(1)} \| P_T^{(0)}) &= -\log(1-q_r) \\ &\le 2q_r \le 2(2+\alpha)^n C e^{-nLT} r^n, \end{aligned}$$

where we use (5) for the equation and the inequality holds when $q_r \le \frac{1}{2}(\frac{2+\alpha}{2})^n \le 0.79$ for sufficiently small α. For N independent endpoint samples,

$$D_{\mathrm{KL}}\Big((P_T^{(1)})^N \| (P_T^{(0)})^N\Big) = N D_{\mathrm{KL}}(P_T^{(1)} \| P_T^{(0)}) \le 2(2+\alpha)^n N C e^{-nLT} r^n. \tag{6}$$

We now relate Hausdorff estimation to binary testing. Define the test induced by $\widehat{S}_N$ as

$$\psi(Y_1, \dots, Y_N) := \begin{cases} 0, & d_H(\widehat{S}_N, \overline{S_T}^{(0)}) \le d_H(\widehat{S}_N, \overline{S_T}^{(1)}), \\ 1, & d_H(\widehat{S}_N, \overline{S_T}^{(1)}) < d_H(\widehat{S}_N, \overline{S_T}^{(0)}). \end{cases}$$

By (4), if $d_H(\widehat{S}_N, \overline{S_T}^{(0)}) \le r$, then $\psi = 0$; similarly, if $d_H(\widehat{S}_N, \overline{S_T}^{(1)}) \le r$, then $\psi = 1$. In other words, the following two implications hold:

$$\begin{cases} d_H(\widehat{S}_N, \overline{S_T}^{(0)}) \le r \implies \psi = 0, \\ d_H(\widehat{S}_N, \overline{S_T}^{(1)}) \le r \implies \psi = 1. \end{cases}$$

Therefore, their contrapositives also hold. Since $\psi \in \{0,1\}$, these contrapositives can be written explicitly as

$$\begin{cases} \psi = 1 \implies d_H(\widehat{S}_N, \overline{S_T}^{(0)}) > r, \\ \psi = 0 \implies d_H(\widehat{S}_N, \overline{S_T}^{(1)}) > r. \end{cases}$$

Equivalently, this gives the event inclusions

$$\begin{cases} \{\psi = 1\} \subseteq \{d_H(\widehat{S}_N, \overline{S_T}^{(0)}) > r\}, \\ \{\psi = 0\} \subseteq \{d_H(\widehat{S}_N, \overline{S_T}^{(1)}) > r\}. \end{cases}$$

Consequently,

$$(P_T^{(0)})^N \left(d_H(\widehat{S}_N, \overline{S_T}^{(0)}) > r \right) \geq (P_T^{(0)})^N(\psi = 1),$$

and

$$(P_T^{(1)})^N \left(d_H(\widehat{S}_N, \overline{S_T}^{(1)}) > r \right) \geq (P_T^{(1)})^N(\psi = 0).$$

Therefore, for this estimator $\widehat{S}_N$,

$$\begin{aligned} \sup_{\ell \in \{0,1\}} (P_T^{(\ell)})^N \left(d_H(\widehat{S}_N, \overline{S_T}^{(\ell)}) > r \right) &\geq \max \left\{ (P_T^{(0)})^N(\psi = 1), (P_T^{(1)})^N(\psi = 0) \right\} \\ &\geq \frac{1}{2} \left[(P_T^{(0)})^N(\psi = 1) + (P_T^{(1)})^N(\psi = 0) \right]. \end{aligned} \tag{7}$$

By Lemma 5,

$$(P_T^{(0)})^N(\psi = 1) + (P_T^{(1)})^N(\psi = 0) \geq \frac{1}{2} \exp\left(-D_{\mathrm{KL}}\left((P_T^{(1)})^N \| (P_T^{(0)})^N \right) \right).$$

Combining this with (7) and (6) gives

$$\sup_{\ell \in \{0,1\}} (P_T^{(\ell)})^N \left(d_H(\widehat{S}_N, \overline{S_T}^{(\ell)}) > r \right) \geq \frac{1}{4} \exp\left(-2(2+\alpha)^n NCe^{-nLT} r^n \right). \tag{8}$$

Now assume

$$N < \frac{e^{nLT}}{2(2+\alpha)^n Cr^n} \log \frac{1}{4\delta}.$$

Then

$$2(2+\alpha)^n NCe^{-nLT} r^n < \log \frac{1}{4\delta}.$$

Substituting this into (8) yields

$$\sup_{\ell \in \{0,1\}} (P_T^{(\ell)})^N \left(d_H(\widehat{S}_N, \overline{S_T}^{(\ell)}) > r \right) > \delta.$$

Since $\widehat{S}_N$ was arbitrary, this means that for every estimator $\widehat{S}_N$, at least one of the two admissible instances $(S_0^{(i)}, F, P_0^{(i)}), i = 1, 2$ satisfies

$$P_T^N \left(d_H(\widehat{S}_N, S_T) > r \right) = P_T^N \left(d_H(\widehat{S}_N, \overline{S_T}) > r \right) > \delta.$$

And such conclusion holds for all sufficiently small α, then we have when

$$N < \lim_{\alpha \to 0^+} \frac{e^{nLT}}{2(2+\alpha)^n Cr^n} \log \frac{1}{4\delta} = \frac{e^{nLT}}{2^{n+1} Cr^n} \log \frac{1}{4\delta}.$$

There always exists an admissible L-Lipschitz dynamic class $F(\cdot)$, with endpoint support $\overline{S_T}$ and endpoint distribution P_T, such that $P_T^N \left(d_H(\widehat{S}_N, \overline{S_T}) > r \right) > \delta$, which proves the minimax lower bound. □

Remark 2 (Worst-case and instance-dependent time dependence)**.** *A natural question is whether the exponential-in-time factor in Theorem 2 is only a pathological worst-case effect. The minimax result is a uniform impossibility statement: it does not imply that every fixed instance requires sample complexity scaling as $e^{nLT} r^{-n}$, where L is a global Lipschitz constant and T is the fixed time*

horizon. For a fixed instance, the relevant quantity is the smallest endpoint mass of an r-scale neighborhood,

$$\gamma_T(r) := \inf_{y \in \overline{S_T}} P_T\big(\mathcal{B}_r(y) \cap \overline{S_T}\big),$$

and the effective sample complexity is governed by $1/\gamma_T(r)$.

The lower-bound construction shows that exponential degradation can occur when the dynamics expands small initial neighborhoods before time T. In the hard instance used in the proof, an endpoint ball of radius r has a preimage with radius on the order of $e^{-LT}r$. Consequently, its initial probability mass is on the order of $e^{-nLT}r^n$, which leads to the lower-bound scaling $N \gtrsim e^{nLT}r^{-n}$ up to constants.

For a general nonlinear system $\dot{x} = F(x)$, the corresponding instance-dependent quantity is controlled by the derivative of the flow map $D_x\varphi(T, x)$. If the flow expands volume over the relevant region, then small endpoint neighborhoods have small preimages under the inverse flow. Writing

$$\Lambda_T := \sup_{x \in S_0} \log |\det D_x\varphi(T, x)|$$

as an effective finite-time volume-expansion exponent over the relevant region, one expects an r-scale endpoint neighborhood to have mass of order $\gamma_T(r) \asymp e^{-\Lambda_T}r^n$,[2] up to density, anisotropy, and curvature-dependent constants. Thus an instance-dependent analysis can replace the worst-case exponent nLT by a smaller effective volume-growth exponent Λ_T, giving the heuristic sample requirement $N \gtrsim e^{\Lambda_T}r^{-n}$. In the isotropic case $\varphi(T, x) \approx e^{\lambda T}x$, we have $\Lambda_T = n\lambda T$, recovering the scaling $e^{n\lambda T}r^{-n}$. Hence, while the minimax theorem uses the worst-case global Lipschitz expansion nLT, persistent positive local volume expansion can still create an exponential-in-time sampling burden for Hausdorff-accurate reachable-set recovery.

C Experiments Details

This section provides additional implementation details and experimental results. We first describe the simulation platform, hardware configuration, robot-arm dynamics, parameter settings, and the estimator based on Christoffel functions used in the supplementary experiments. We then report additional results on dimension-dependent approximation of robot-arm uncertainty propagation, followed by reachable-set approximation results obtained using the Christoffel estimator.

C.1 Experimental Setup

Simulation Platform and Hardware. All simulations are run on a 3.2 GHz AMD Ryzen 7 7735HS CPU with 16 GB RAM. The robot-arm simulations are implemented in MuJoCo 3.8.1. We use deterministic MuJoCo MJCF models with a fixed integration time step of 2×10^{-3} seconds. The approximation error is computed offline from simulated terminal samples. In the main experiments in Section 5, we use the convex hull of the sampled endpoints as the reachable-set estimator.

Dynamics and Parameter Settings of Robotic Arms. We consider vertical planar serial n-link robot arms with $n \in \{2, 3, 4\}$. The state is $x = [q^\top, v^\top]^\top \in \mathbb{R}^{2n}$, where $q \in \mathbb{R}^n$ denotes joint angles and $v = \dot{q} \in \mathbb{R}^n$ denotes joint velocities. Each link has length 0.5, capsule radius 0.035, density 1000, joint damping 0.2, and armature 0.01. Gravity is enabled with acceleration $(0, 0, -9.81)$. MuJoCo simulates the rigid-body dynamics

$$M(q)\dot{v} + C(q, v)v + g(q) = \tau,$$

where $\tau \in \mathbb{R}^n$ is the vector of joint torques. For the simulation, we consider a non-adaptive inverse-dynamics tracking controller. The desired reference trajectory is

$$q_{d,i}(t) = q_{c,i} + A_i \sin(\omega t + \phi_i),$$

[2] $A \lesssim B$ means $A \leq cB$, $A \gtrsim B$ means $A \geq cB$, and $A \asymp B$ means both hold, for constants independent of the scale parameter.

with $q_c = \mathrm{linspace}(0.25, 0.65, n)$, $A_i = 0.08$, $\omega = 0.5$, $\phi = \mathrm{linspace}(0, \pi/3, n)$. The analytic derivatives $\dot{q}_d(t)$ and $\ddot{q}_d(t)$ are used in the controller. Let $e = q - q_d(t)$ and $\dot{e} = v - \dot{q}_d(t)$. MuJoCo inverse dynamics is used to compute

$$\tau_{\mathrm{track}} = M(q)\ddot{q}_d(t) + C(q, v)v + g(q).$$

The applied torque is

$$\tau = \tau_{\mathrm{track}} - K_p e - K_d \dot{e}.$$

We use weak tracking gains

$$K_p = 0.01 I_n, \qquad K_d = 0.005 I_n,$$

and clip torques to $[-100, 100]$ for numerical stability. The weak gains are chosen deliberately: the benchmark is designed for uncertainty propagation and finite-sample reachable-set approximation, rather than aggressive trajectory tracking or controller design.

The initial uncertainty set is the box $S_0 = ([-\rho_q, \rho_q]^n \times [-\rho_v, \rho_v]^n) \circ \mathcal{B}(0, r)$ with $\rho_q = \rho_v = 0.1, r = 0.01$. Each trajectory is propagated to time $T = 1.0$ second, and the terminal propagated uncertainty set is approximated from sampled initial conditions.

Estimator using Christoffel function. To further validate the geometric implications of our theory, we use the empirical inverse Christoffel function as an independent verification tool [17]. Given endpoint samples $\{Y_i\}_{i=1}^N$, we construct

$$\widehat{M}_{m,\sigma_0} = \sigma_0^2 I + \frac{1}{N} \sum_{i=1}^{N} z_m(Y_i) z_m(Y_i)^\top,$$

where z_m is the vector of monomials of degree at most m, and define

$$C(y) = z_m(y)^\top \widehat{M}_{m,\sigma_0}^{-1} z_m(y).$$

Following the polynomial Christoffel estimator in [18, Algorithm 3], we use the sublevel set $\widehat{S}_N^{\mathrm{Ch}} = \{y : C(y) \le \eta\}$ as a data-driven certificate of support coverage. In our setting, this construction is not used as a competing reachable-set estimator, but rather as an auxiliary verification procedure: if the endpoint samples generated under our sampling scheme are sufficiently informative, the resulting Christoffel sublevel set should concentrate around the reachable set and provide an independent check of the predicted finite-sample behavior.

C.2 Adversarial Sampling Details

For completeness, we describe the adversarial sampling heuristic over S_0 used in the experiments to improve empirical coverage [61, 62]. Given the current endpoint cloud $\mathcal{Y}^i$, define $c^i := |\mathcal{Y}^i|^{-1} \sum_{Y \in \mathcal{Y}^i} Y$ and $Q^i := [\frac{1}{|\mathcal{Y}^i|-1} \sum_{Y \in \mathcal{Y}^i} (Y - c^i)(Y - c^i)^\top + \lambda I]^{-1}$, where $\lambda > 0$ is a regularization parameter. We then maximize the novelty objective $\mathcal{L}^i(x) := \|\varphi(T, x) - c^i\|_{Q^i}^2$,

Algorithm 1 Reachable Set Estimation via Adversarial Sampling

Input: Initial samples $\{X_j^0\}_{j=1}^M \subset S_0$
Parameters: Stepsize η, iterations n_{adv}
Output: Estimate $\widehat{S}_N^{\mathrm{adv}}$
$Y_j^0 \leftarrow \varphi(T, X_j^0), \quad j = 1, \dots, M$
$\mathcal{Y}^0 \leftarrow \{Y_j^0\}_{j=1}^M$
for $i = 0, \dots, n_{\mathrm{adv}} - 1$ **do**
    Compute c^i, Q^i from $\mathcal{Y}^i$
    for $j = 1, \dots, M$ **do**
        $X_j^{i+1} \leftarrow \mathrm{Proj}_{S_0}\big(X_j^i + \eta \nabla_x \mathcal{L}^i(X_j^i)\big)$
        $Y_j^{i+1} \leftarrow \varphi(T, X_j^{i+1})$
    end for
    $\mathcal{Y}^{i+1} \leftarrow \mathcal{Y}^i \cup \{Y_j^{i+1}\}_{j=1}^M$
end for
$N \leftarrow M(n_{\mathrm{adv}} + 1)$
return $\widehat{S}_N^{\mathrm{adv}} = \mathrm{C}(\mathcal{Y}^{n_{\mathrm{adv}}})$

$x \in S_0$, which encourages new samples to generate endpoints far from the current cloud after covariance normalization [19]. Thus, the procedure allocates samples toward geometrically underexplored regions of the reachable set. At each iteration, the current endpoint cloud is updated along $\nabla_x \mathcal{L}^i$, projected back to S_0, and propagated through the dynamics. The newly generated endpoints are then added to the sample cloud, and the final estimator is applied to all accumulated endpoints. For more details, please refer to [19].

C.3 Additional Results

Fitted Dimension-Dependent Slopes. To summarize the dimension-dependent trend in Table 2, we fit the magnitude of the empirical log–log slope as a function of the state dimension n using

$$m(n) = \frac{1}{an^b + c}, \qquad \text{slope}(n) \approx -m(n).$$

Since the fitted curves are obtained from plots of $\log d_H$ versus $\log N$, they can also be interpreted as empirical sample-complexity exponents. Indeed, if

$$\log d_H \approx \alpha(n) - m(n) \log N,$$

then $d_H \approx A_n N^{-m(n)}$, where $A_n = e^{\alpha(n)}$ is the dimension-dependent prefactor in the empirical scaling law; equivalently, it is the value predicted by the fitted power law at $N = 1$. In practice, A_n should be interpreted as an intercept parameter rather than as a reliable one-sample approximation error, since the fit is obtained over a finite range of sample sizes. This prefactor captures effects not represented by the slope alone, including the geometric scale of the reachable set, flow-induced expansion, sampling density, and estimator bias. Empirically, A_n often increases with the state dimension, which further amplifies the sample requirement. This interpretation is consistent with classical support- and level-set estimation results, where Hausdorff-type recovery rates and covering complexities depend explicitly on the ambient or intrinsic dimension [63]. Thus achieving accuracy $d_H \le r$ requires, up to dimension-dependent constants, $N \gtrsim \left(\frac{A_n}{r}\right)^{1/m(n)} = \left(\frac{A_n}{r}\right)^{an^b+c}$. Therefore, the fitted denominator $an^b + c$ represents the empirical exponent with which the sample budget must grow as the target accuracy r decreases.

The corresponding parameters and fitted curves are shown in Table 3 and Figure 7, respectively.

Substituting the fitted parameters in Table 3 into the relation $N \gtrsim \left(\frac{A_n}{r}\right)^{an^b+c}$ shows that the sample budget grows as

$$N_{\mathrm{unif}} \gtrsim \left(\frac{A_n}{r}\right)^{0.5049\, n^{1.0709}+1.3353}, N_{\mathrm{adv}} \gtrsim \left(\frac{A_n}{r}\right)^{0.9488\, n^{0.7203}+0.2535}.$$

Table 3: Fitted parameters for the empirical slope model $m(n) = 1/(an^b + c)$, where n is the state dimension and slope$(n) \approx -m(n)$. The implied sample-complexity exponent is $1/m(n) = an^b + c$.

Sampling method	a	b	c
Uniform sampling	0.5049	1.0709	1.3353
Adversarial sampling	0.9488	0.7203	0.2535

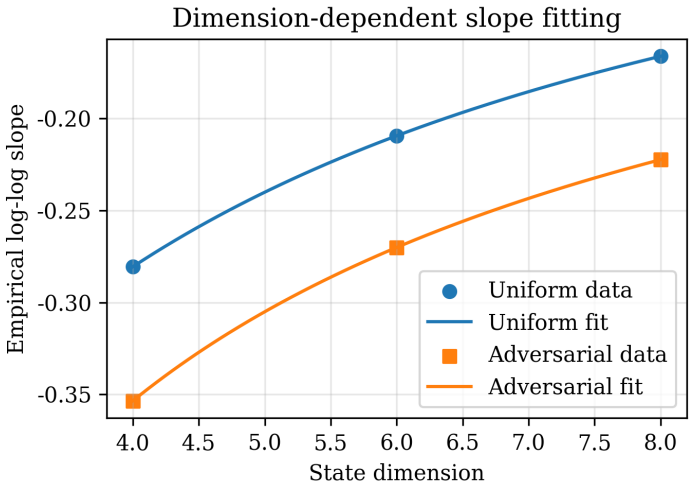

Figure 7: Fitted dimension-dependent empirical slopes for robot-arm uncertainty propagation. The fitted model captures the flattening of the log–log convergence slopes as the state dimension increases.

Equivalently, taking logarithms on both sides yields

$$\log N_{\text{unif}} \gtrsim \left(0.5049\, n^{1.0709} + 1.3353\right) \log \frac{A_n}{r}, \qquad \log N_{\text{adv}} \gtrsim \left(0.9488\, n^{0.7203} + 0.2535\right) \log \frac{A_n}{r}.$$

Thus the empirical sample requirement does not grow only logarithmically with dimension. Both uniform and adversarial sampling exhibit an exponential-type dependence on the state dimension through the exponent multiplying $\log(A_n/r)$. Adversarial sampling reduces the fitted exponent, but the exponent still increases with n, indicating that targeted sampling improves finite-sample efficiency without eliminating the intrinsic curse of dimensionality.

Sample-Budget-Dependent Improvement of Adversarial Sampling. Section 5 reports the main dimension-dependent robot-arm results. Here we provide an additional comparison showing how the benefit of adversarial sampling depends on the sample budget. Table 4 reports the absolute and relative improvement of adversarial sampling over uniform sampling for the 2-, 3-, and 4-link robot arms.

The advantage of adversarial sampling is not uniform across all sample sizes. When N is very small, adversarial updates may over-concentrate samples along a few extreme directions and reduce global coverage. As the sample budget grows, the sampler has enough coverage to exploit informative directions, and the relative improvement becomes positive and more stable. At $N = 3000$, the relative improvements are approximately 39.09%, 35.56%, and 30.58% for the 2-, 3-, and 4-link arms, respectively.

Time dependence under closed-loop control. Figure 8 shows that, under the closed-loop robot-arm dynamics, the Hausdorff error grows only slowly with the simulation horizon T. This contrasts with the non-Lipschitz and Linear system examples in Section 5.1, where flow expansion strongly amplifies sampling errors. In the robot-arm setting, the inverse-dynamics tracking controller reduce

Table 4: Absolute (Abs.) and relative (Rel.) improvement of adversarial sampling over uniform sampling for robot-arm uncertainty propagation.

	$n = 2$		$n = 3$		$n = 4$	
N	Abs.	Rel. (%)	Abs.	Rel. (%)	Abs.	Rel. (%)
1	-0.0117	-3.35	0.0000	0.01	-0.0098	-2.15
3	0.0236	8.27	-0.0157	-4.88	-0.0126	-3.47
10	0.0356	17.79	0.0122	4.93	0.0011	0.37
30	0.0365	25.86	0.0374	18.01	0.0250	10.26
100	0.0397	36.72	0.0316	20.70	0.0386	18.84
300	0.0317	40.10	0.0307	25.35	0.0462	25.90
1000	0.0226	41.17	0.0290	30.12	0.0371	26.68
3000	0.0144	39.09	0.0280	35.56	0.0355	30.58

the effective expansion of the uncertainty set, so the observed time dependence is much milder than the worst-case exponential growth predicted by the minimax bound.

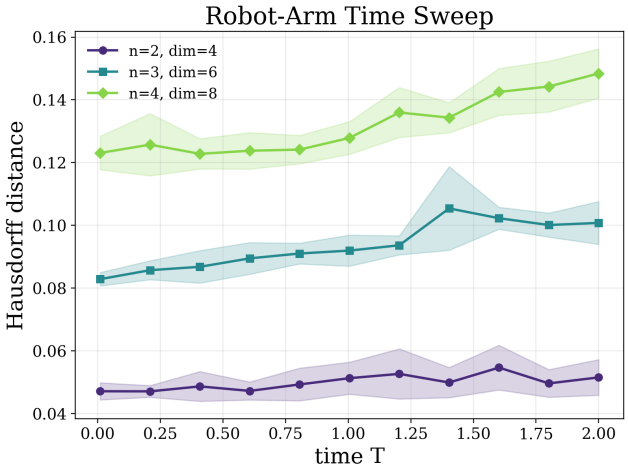

Figure 8: Hausdorff error versus time under uniform sampling

Effect of Adversarial Sampling Intensity Across Estimators. We further study how the adversarial intensity of the sampling procedure affects reachable-set approximation under different downstream estimators. We repeat the time-sweep experiment for the autonomous system $\dot{y} = 0, \dot{x} = x^2$ using the same three initial sets as in Section 5.1: a disk, a triangle, and an opened triangle with equal area. We compare two estimators constructed from the same endpoint samples: the convex hull estimator and a Christoffel-type estimator.

The adversarial intensity is controlled by the number of adversarial updates $n_{\text{adv}} \in \{0, 1, 2, 3, 4\}$, where $n_{\text{adv}} = 0$ corresponds to purely uniform sampling. Larger n_{adv} allocates a larger fraction of the fixed sample budget to points obtained after adversarial updates, and therefore places more emphasis on expanding or boundary-like directions in the endpoint cloud.

Figures 9 and 10 show the Hausdorff error over time for sample budgets $N = 10, 100, 1000$. The results show that the effect of adversarial sampling is both budget-dependent and estimator-dependent. For very small sample budgets, increasing n_{adv} does not consistently improve performance, since overly concentrated adversarial samples may reduce global coverage. For larger budgets, adversarial updates become more beneficial: the endpoint cloud already has enough global coverage, and additional adversarial samples help resolve boundary and extreme regions. This trend is visible for both convex-hull and Christoffel estimators, although the best choice of n_{adv} varies across estimators, sample budgets, and initial geometries. Overall, these results suggest a coverage–boundary tradeoff:

stronger adversarial sampling can improve finite-sample performance, but only when the sample budget is large enough to avoid sacrificing global coverage.

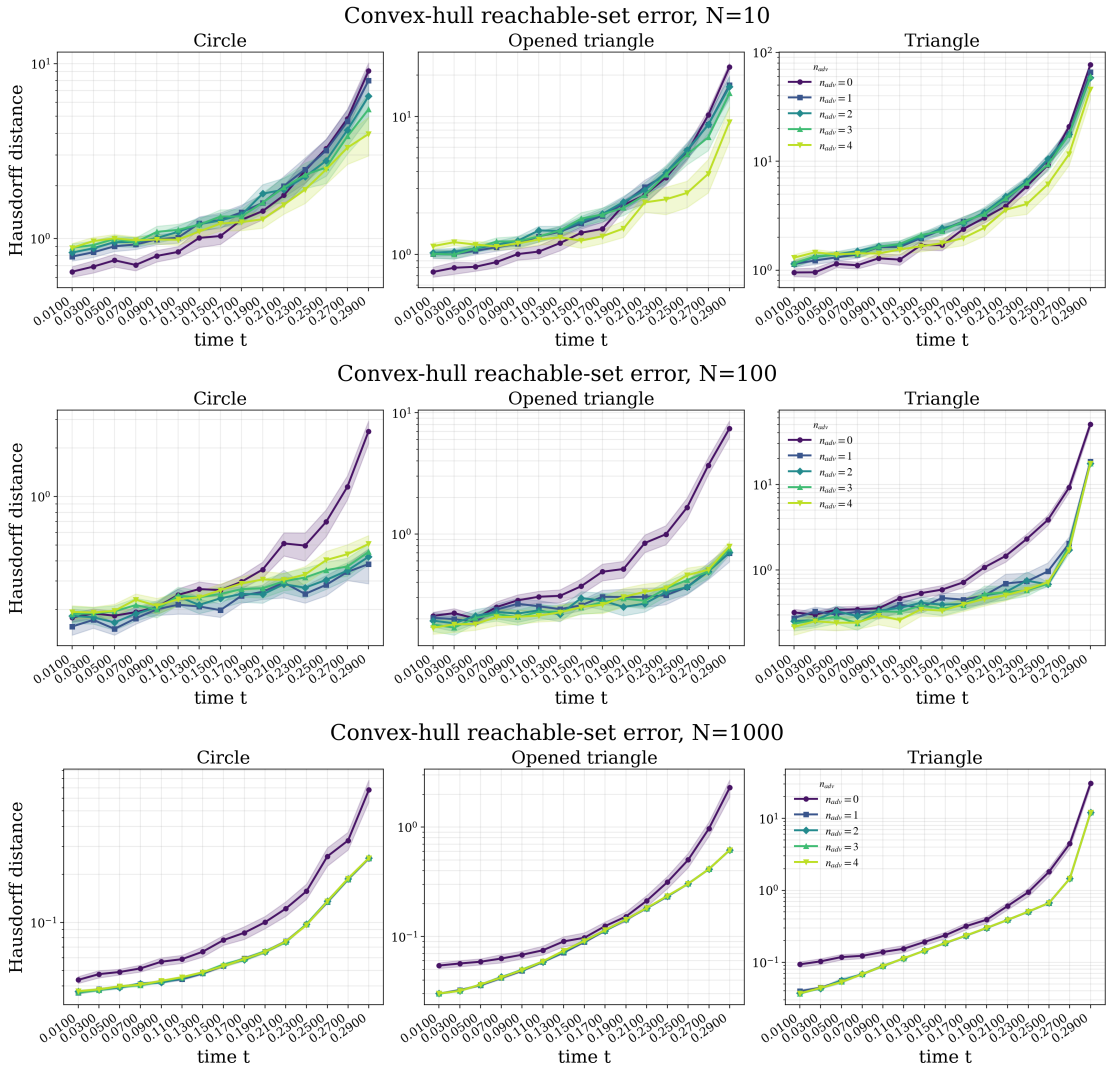

Figure 9: Hausdorff error versus time for Convexhull reachable-set approximation of the autonomous system $\dot{y} = 0, \dot{x} = x^2$ under different uniform–adversarial mixture sampling ratios. The three subfigures correspond to sample budgets $N = 10, 100, 1000$, and each subfigure compares the disk, triangle, and opened-triangle initial sets.

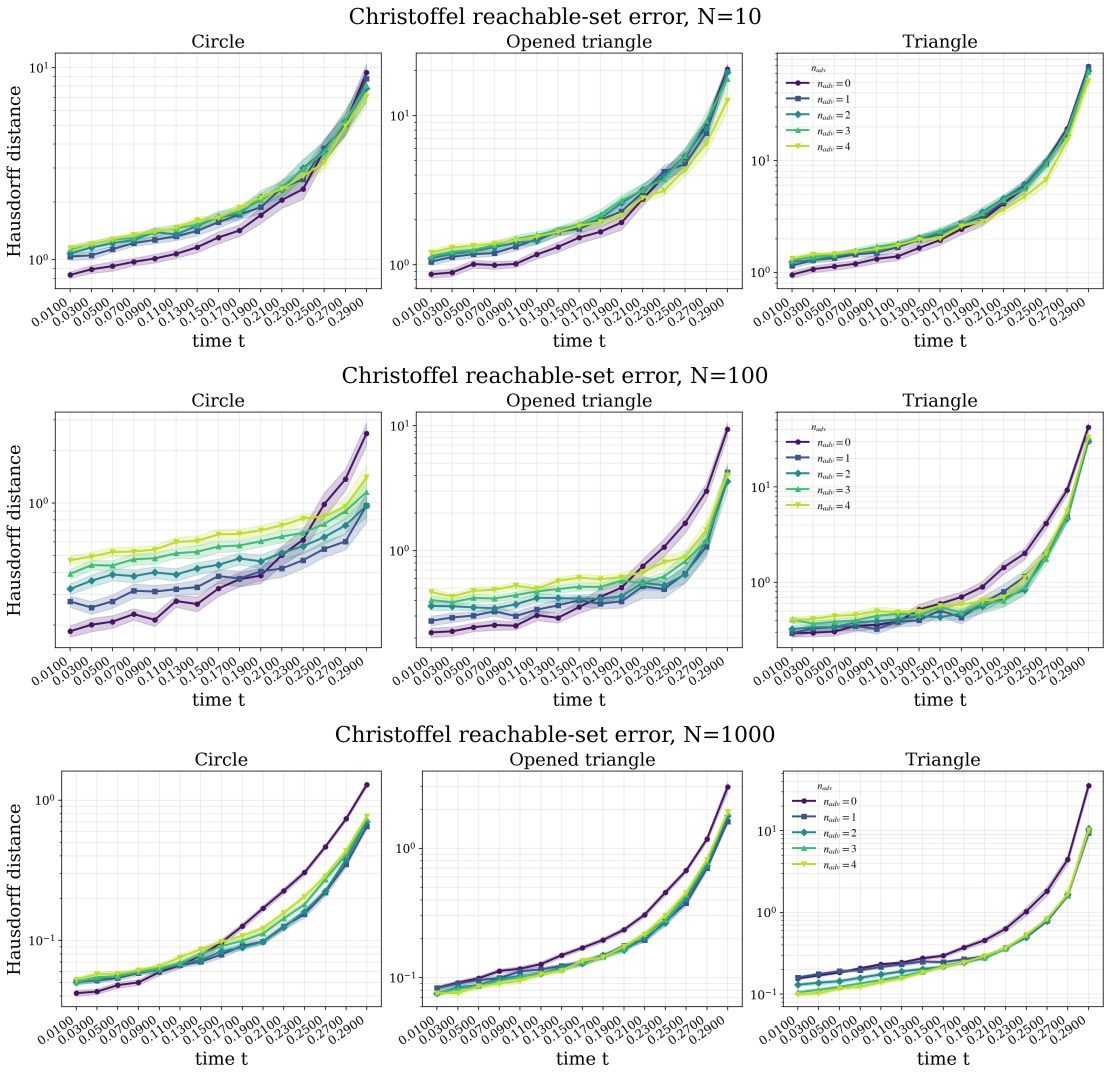

Figure 10: Hausdorff error versus time for Christoffel reachable-set approximation of the autonomous system $\dot{y} = 0, \dot{x} = x^2$ under different uniform–adversarial mixture sampling ratios. The three subfigures correspond to sample budgets $N = 10, 100, 1000$, and each subfigure compares the disk, triangle, and opened-triangle initial sets.