

Manifold-Guided Lyapunov Control with Diffusion Models

Amartya Mukherjee, Thanin Quartz, and Jun Liu*

Abstract

This paper presents a novel approach to generating stabilizing controllers for a large class of dynamical systems using diffusion models. The core objective is to develop stabilizing control functions by identifying the closest asymptotically stable vector field relative to a predetermined manifold and adjusting the control function based on this finding. To achieve this, we employ a diffusion model trained on pairs consisting of asymptotically stable vector fields and their corresponding Lyapunov functions. Our numerical results demonstrate that this pre-trained model can achieve stabilization over previously unseen systems efficiently and rapidly, showcasing the potential of our approach in fast zero-shot control and generalizability.

1 INTRODUCTION

One of the longstanding challenges in nonlinear systems and control is the construction of Lyapunov functions. While Lyapunov functions can essentially be characterized by solutions to partial differential equations (PDEs) and neural network solutions to such PDEs can effectively provide Lyapunov functions [12], solving PDEs for each system can still be time-consuming. Another challenge in computing neural Lyapunov functions is the inclusion of a satisfiability modulo theories (SMT) solver for counterexample-guided training. Verifying that a neural network satisfies all the conditions necessary to be a Lyapunov function is often computationally expensive and significantly increases the training time of such methods. Additionally, the neural Lyapunov function only aids in the verification of a single control system, thus failing to take into account possible uncertainties in the model parameters or the ability to generalize into systems with slightly different dynamics. This could pose difficulties when testing the controller in real-life systems.

Generative models, such as diffusion models [8, 20], map inputs from a noisy space to the data space. In more recent works, diffusion models have learned to transform noisy inputs or boundary conditions into solutions for PDEs, including the Poisson equation and the Navier-Stokes equation, without incorporating any information about the PDEs into their loss function [2, 15]. Furthermore, denoising diffusion restoration models (DDRM) [10, 14] can restore clean data in linear inverse problems using a pre-trained diffusion model, eliminating the need for fine-tuning.

In this paper, we investigate the potential of using pre-trained generative diffusion models to generate Lyapunov functions and stabilizing controllers for nonlinear systems. More specifically, we aim to train a diffusion model, G , that generates pairs of asymptotically stable vector fields and their corresponding Lyapunov functions. This simplifies the control problem to finding a controller such that the closed-loop system possesses a Lyapunov-stable vector field, indicating it falls within the range of G . The Lyapunov function produced by G further assists in verifying stability for the control problem. This approach offers several advantages over existing learning-based methods for Lyapunov functions (e.g., [3, 23]): (a) it can significantly reduce the computational effort required to identify a stabilizing controller, (b) it has the capability to generalize to control problems not encountered in the training data, and (c) it can offer stability guarantees with the generated Lyapunov functions.

We propose **Manifold Guided Lyapunov Control (MGLC)**, a novel approach to feedback control problems based on diffusion models. We assume that the stabilized controlled system is drawn from a manifold \mathcal{M} that contains asymptotically stable dynamical systems. We also introduce a manifold \mathcal{W} that contains all possible configurations of the dynamical system with different control functions. This simplifies the control problem to finding points where \mathcal{M} and \mathcal{W} intersect. We train a diffusion model with pairs of asymptotically stable vector fields and their Lyapunov functions to estimate \mathcal{M} and incorporate the control design process into the reverse diffusion process. At every time step, we estimate the projection of our vector field into \mathcal{M} by computing Tweedie’s estimate, and we update the parameters of the controller to minimize a loss function that reduces the distance between our vector field between the two vector fields.

*Amartya Mukherjee, Thanin Quartz, and Jun Liu are with the Department of Applied Mathematics, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1 (email: (a29mukhe, tquartz, j.liu)@uwaterloo.ca).

Our proposed solution to this problem significantly reduces the computational cost of finding a Lyapunov function. Rather than training a neural network for a specific system, often facilitated by an SMT solver for generating counterexamples [1], we use a pre-trained diffusion model to directly output a Lyapunov function that verifies the convergence of our system. As a result, this approach will be useful for solving a large class of control problems rather than a single control problem, thus facilitating zero-shot control, where we achieve stabilization of unseen systems. We introduce theoretical guarantees that our reverse diffusion sampling method will converge to the manifold \mathcal{M} , thus concluding that we have derived a stabilizing controller. We train the generator on a dataset of pairs of asymptotically stable vector fields and their Lyapunov functions. Finally, we test our method on four control problems whose dynamics are not in the dataset and are therefore unseen.

Notation

Denote X as the discretization of the domain into a grid and $f(X, u(X))$ as the vector field f projected onto the grid X with controller $u(\cdot)$. We use $\psi = (\psi_1, \dots, \psi_n)$ to refer to the learnable parameters for a controller $u^{(\psi)}(\cdot)$, where n is the number of learnable parameters. Subsequently, ψ_t are the parameters of the control system at time step t in the denoising diffusion process. Therefore, $x_t := f(X, u^{(\psi_t)}(X))$ denotes the projection of the vector field $f(\cdot, u^{(\psi_t)}(\cdot))$ on the gridpoints X at time step t in the denoising diffusion process, where the parameters of the controller are ψ_t . Lastly, $x_{0|t} := E[x_0|x_t]$ refers to the estimate of the asymptotically stable vector field x_0 conditioned on the vector field at time step t , x_t . This estimate is computed using Tweedie’s formula, which we will formally define in Section 2.2.

2 PRELIMINARIES AND PROBLEM FORMULATION

2.1 Lyapunov Stability

Throughout this work we consider a control system of the following form

$$\dot{x} = f(x, u), \quad x(0) = x_0, \quad (1)$$

where $x \in D$ is the state space of the system, $D \subset \mathbb{R}^n$ is an open set containing the origin and $u \in U \subset \mathbb{R}^m$ is the feedback control input given by $u = u(x)$, where $u(\cdot)$ is a continuous function of the state. Moreover, we assume that f is a locally Lipschitz continuous vector field, $u(\cdot)$ is locally Lipschitz with $u(0) = 0$, and $f(0, 0) = 0$.

Definition 1. (Asymptotic Stability) *The system (1), under a feedback controller $u = u(x)$, is asymptotically stable at the origin if for any $\varepsilon \in \mathbb{R}^+$, there exists $\delta(\varepsilon) \in \mathbb{R}^+$ such that if $\|x_0\| < \delta$, then $\|x(t)\| < \varepsilon$ for all $t \geq 0$. Furthermore, there exists some $\rho > 0$ such that $\|x_0\| < \rho$ implies $\lim_{t \rightarrow \infty} \|x(t)\| = 0$.*

Definition 2 (Lie Derivatives). *Let $D \subset \mathbb{R}^n$. The Lie derivative of a continuously differentiable scalar function $V : D \rightarrow \mathbb{R}$ over a vector field $f(\cdot, u(\cdot))$ and feedback controller $u(\cdot)$ is defined as*

$$\dot{V}(x) = \sum_{i=1}^n \frac{\partial V}{\partial x_i} \dot{x}_i = \sum_{i=1}^n \frac{\partial V}{\partial x_i} f_i(x, u(x)) \quad (2)$$

Theorem 1. *If there exists a continuously differentiable scalar function $V(x)$ satisfying $V(0) = 0$, $V(x) > 0$ if $x \neq 0$, and $\nabla V(x) \cdot f(x, u(x)) < 0$, for a given feedback controller u , then the system is asymptotically stable (or Lyapunov-stable) at the origin, and V is called a Lyapunov function [3].*

Consequently, in control problems $\dot{x} = f(x, u)$, a feedback control function $u = k(x)$ is a stabilizing controller if $\dot{x} = f(x, k(x))$ is an asymptotically stable system. The existence of a Lyapunov function verifies the stability of the system. However, verifying the Lyapunov conditions is a difficult task that can be handled by SMT solvers as in [3, 23] to identify regions where the neural network V violates the conditions for it to be a Lyapunov function and a penalty term for those regions is added to its loss function.

2.2 Diffusion Models

Denoising diffusion models [19, 20] are among the current leading methods for generative modeling [8, 18]. They have shown great success in applications such as the generation of images, speech, and video, as well as image super-resolution [18, 21]. Other applications include physics-guided human motion (e.g., PhysDiff [22]), customized ODE solvers that are more efficient than Runge-Kutta methods [13], molecule generation [9], and more [21]. Furthermore, they are stable to train and are relatively easy to scale [21].

Diffusion models consist of a forward process where data is iteratively corrupted by adding noise. This is modeled by the following process:

$$x_t = \sqrt{\bar{\alpha}_t}x + \sqrt{1 - \bar{\alpha}_t}\varepsilon_t, \quad t \in [0, T], \quad (3)$$

where $\bar{\alpha}_t > 0$ is a scaling parameter that monotonically decreases with t , $\varepsilon_t \sim \mathcal{N}(0, I)$ is Gaussian noise, and $x \sim p_{\text{data}}$ is clean data. The goal of a diffusion model is to reverse this process. This is done by introducing a denoising model $\varepsilon_\theta(x_t, t)$ trained with the loss function $\|\varepsilon_t - \varepsilon_\theta(x_t, t)\|^2$ to estimate the Gaussian noise. This model approximates the score function $\nabla_{x_t} \log p_t(x_t)$.

This methodology allows us to generate realistic data. Particularly, denoising diffusion implicit models (DDIM) [18] perform sampling using a two-step procedure for multiple time steps. The first step involves computing $x_{0|t} := E[x_0|x_t]$ using Tweedie's estimate, which we will define below.

Definition 3 (Tweedie's estimate [6, 5]). *The Tweedie's estimate for $E[x_0|x_t]$ governed by the process in equation (3) is given by the following expression:*

$$x_{0|t} = \frac{x_t - \sqrt{1 - \bar{\alpha}_t}\varepsilon_\theta(x_t, t)}{\sqrt{\bar{\alpha}_t}}. \quad (4)$$

The DDIM algorithm consists of the estimate of $x_{0|t}$ in equation (4) coupled with the estimate of x_{t-1} by the following process:

$$x_{t-1} = \sqrt{\bar{\alpha}_{t-1}}x_{0|t} + \sqrt{1 - \bar{\alpha}_{t-1}}\varepsilon_\theta(x_t, t). \quad (5)$$

This process is repeated for multiple time steps. Note as we denoted previously, x_t denotes the discretization of the right-hand side of (1) on grid points X at time t in the denoising diffusion process. Therefore, (5) says we are updating discretized values of the right hand corresponding to updated parameters of the controller.

2.3 Diffusion on Manifolds

Diffusion models for inverse problems have gained interest with the advent of denoising diffusion restoration models (DDRM) [10], originally used to deblur images by assuming that the original clean image was drawn from a pre-trained unconditional diffusion model. This field, otherwise known as guided diffusion, has been improved upon with diffusion posterior sampling [4], partially collapsed Gibbs sampler [14], and manifold constraints [4].

The problem of guided diffusion involves sampling from a conditional score function $\nabla_{x_t} \log p(x_t|y)$ conditioned on a variable y instead of sampling from the data distribution $\nabla_{x_t} \log p(x_t)$ as done in the pre-trained diffusion model. In many applications, this conditioning y can be a text input [16], an image input, or an image that we intend to restore [14]. This conditional score function can be separated into two terms using Bayes' theorem: $\nabla_{x_t} \log p(x_t|y) = \nabla_{x_t} \log p(x_t) + \nabla_{x_t} L_t(x_t; y)$. Sampling from $\nabla_{x_t} \log p(x_t)$ is done by our diffusion model, but the sampling procedure $\nabla_{x_t} L_t(x_t; y)$ must be carefully chosen.

The authors of [4] assume that the data distribution has support in a low-dimensional manifold, known as the manifold hypothesis. They exploit this property to enable conditional sampling on an unconditionally trained diffusion model, thus achieving state-of-the-art performance in inpainting, colorization, and reconstruction tasks in images.

The work of [7] builds upon this literature by using tangent spaces on the data manifold \mathcal{M} to return improved estimates of the original clean data. This method shows state-of-the-art results in image quality and image restoration speed.

2.4 Problem Formulation

In this paper, we will consider nonlinear control systems of the following form

$$\dot{x} = f(x, u), \quad x(0) = x_0, \quad (6)$$

where $x \in D$ is the state of the system and $u(\cdot)$ is the feedback control input. Our goal is to derive a controller $u(\cdot)$ such that the trajectory of the system converges to the origin for all $x_0 \in D$.

3 MANIFOLD-GUIDED LYAPUNOV CONTROL (MGLC)

Define the manifold \mathcal{M} of asymptotically stable control vector fields as

$$\mathcal{M} = \{f(X, u(X)) : u \text{ is a stabilizing state feedback controller}\}. \quad (7)$$




Figure 1: A schematic view of MGLC. Our proposed method introduces a new manifold \mathcal{W} corresponding to the set of vector fields that can be returned by modifying our control input function.

We build upon past works in diffusion on manifolds by constructing a manifold $\mathcal{W} = \{f(X, u^{(\psi)}(X)) : \psi \in \Psi\}$ which is the set of vector fields with our parameterized control input function $u^{(\psi)}(\cdot)$. Therefore, we want to learn a controller $u^{(\psi)}$ such that $f(X, u^{(\psi)}(X)) \in \mathcal{M} \cap \mathcal{W}$. A schematic view is posted in Figure 1, comparing DDIM [18], MPGD [7], and our proposed method (MGLC).

3.1 Proposed Method

We introduce a manifold hypothesis for our problem that assumes our dataset of interest, that is discretized values of the set of asymptotically stable vector fields, lies on a manifold that is low-dimensional compared to the set of all functions defined in a domain containing X projected onto X .

Assumption 1 (Manifold Hypothesis). *The support \mathcal{X} of the set of asymptotically stable vector fields projected onto a grid X lies on a k -dimensional manifold $\mathcal{M} \subset \mathbb{R}^d$ with $k \ll d$.*

Compared to [7], we introduce a new manifold \mathcal{W} , which constrains the vector fields our model can output to the set of vector fields with all possible configurations of parameters for our control function.

Assumption 2. *The manifold $\mathcal{W} = \{f(X, u^{(\psi)}(X)) : \psi \in \Psi\}$ contains the vector field f with all possible control parameter configurations. We assume $\mathcal{W} \cap \mathcal{M}_t \neq \emptyset$ for all $t = 0, \dots, T$, in other words, \mathcal{W} intersects every manifold \mathcal{M}_t .*

We also assume that the manifolds \mathcal{M}_t and \mathcal{M}_{t-1} are "close enough" so that a path along the tangent of \mathcal{W} from $\mathcal{M}_t \cap \mathcal{W}$ intersects \mathcal{M}_{t-1} close to a point in $\mathcal{M}_{t-1} \cap \mathcal{W}$.

Assumption 3. *For any time step $t = 1, \dots, T$, the manifolds \mathcal{M}_t and \mathcal{M}_{t-1} are "close enough" with respect to the manifold \mathcal{W} such that a first-order path along \mathcal{W} is a good approximation. Rigorously, for every point $x_t \in \mathcal{M}_t \cap \mathcal{W}$, there exists a smooth path $\gamma : [0, 1] \rightarrow \mathcal{W}$ such that $\gamma(0) = x_t$ and $\gamma(1) = y$, where y is a point in \mathcal{M}_{t-1} satisfying two properties:*

1. *The path γ is entirely contained in the tangent space $T_{x_t}\mathcal{W}$. This means y is contained in $T_{x_t}\mathcal{W} \cap \mathcal{M}_{t-1}$.*
2. *The Euclidean distance $d(y, \mathcal{M}_{t-1} \cap \mathcal{W})$ is less than ε for an arbitrarily small $\varepsilon > 0$.*

We then define a notion of probabilistic concentration of noisy data on a manifold to aid in our convergence results in this paper. This was introduced in the work of [5] and then improvised in the paper by [7] to fit with the discrete-time setting of diffusion models.

Proposition 1 (Probabilistic Concentration [5, 7]). *The distribution of noisy data $p_i(x_i) = \int p(x_i|x)p_0(x)dx$ with $p(x_i|x) \sim \mathcal{N}(\sqrt{\bar{\alpha}_t}x, (1 - \bar{\alpha}_t)I)$ is concentrated on the $(d-1)$ -dimensional manifold $\mathcal{M}_t := \{y \in \mathbb{R}^n : d(y, \sqrt{\bar{\alpha}_t}M) = \sqrt{(1 - \bar{\alpha}_t)(d-k)}\}$. Rigorously, for all $\varepsilon > 0$, there exists a $\delta > 0$ such that $p_t(B_{\varepsilon\sqrt{(1-\bar{\alpha}_t)(d-k)}}(\mathcal{M}_t)) > 1 - \delta$.*

We compute the estimate of x_0 conditioned on x_t , $x_{0|t} := E[x_0|x_t]$ by computing the Tweedie's estimate

$$x_{0|t} = \frac{1}{\sqrt{\bar{\alpha}_t}}(x_t - \sqrt{1 - \bar{\alpha}_t}\varepsilon_\theta(x_t, t)) \quad (8)$$

We then update the parameters ψ using the following gradient update

$$\psi_{t-1} = \psi_t - c_t \nabla_\psi L_t(x_t; \psi), \quad (9)$$

with some loss function $L_t(x_t; \psi)$. Note that Tweedie's estimate is used in the computation of the loss function as described in Algorithm 1. Finally after updating the parameters ψ , we can estimate x_{t-1} :

$$x_{t-1} = f(X, u^{\psi_{t-1}}(X)). \quad (10)$$

3.2 Convergence Proofs

We introduce the following theorems that highlight our main theoretical contribution in this paper. They compute the loss function $L_t(x_t; \psi_t)$ and step size c_t and provides convergence guarantees in the sense that the vector field $f(X, u^{\psi_0}(X))$ returned by our proposed methodology has marginal distribution concentrated about $\mathcal{W} \cap \mathcal{M}$ and therefore $u^{\psi_0}(X)$ is a stabilizing controller for the system. We first introduce a lemma that explains the choice of $L_t(x_t; \psi)$ and c_t .

Proposition 2. *Fix t and assume $f(X, u^{\psi_t}(X)) \in \mathcal{M}_t$. Then with loss function*

$$\begin{aligned} L_t(x_t; \psi) &= \left\| f(X, u^{(\psi_t)}(X)) - \frac{\sqrt{\bar{\alpha}_{t-1}} - \frac{\sqrt{1 - \bar{\alpha}_{t-1}}\sqrt{\bar{\alpha}_t}}{\sqrt{1 - \bar{\alpha}_t}}}{1 - \frac{\sqrt{1 - \bar{\alpha}_{t-1}}}{\sqrt{1 - \bar{\alpha}_t}}} x_{0|t} \right\|^2 \\ &=: \|r_t\|^2, \end{aligned} \quad (11)$$

where we assume that r_t lies on the tangent space of \mathcal{W} at the point $f(X, u^{\psi_t}(X))$, and step size

$$c_t = (1 - \frac{\sqrt{1 - \bar{\alpha}_{t-1}}}{\sqrt{1 - \bar{\alpha}_t}}) \quad (12)$$

we have that there exists x_{t-1} probabilistically concentrated in \mathcal{M}_{t-1} such that

$$x_{t-1} = f(X, u^{\psi_t}(X)) - c_t r_t. \quad (13)$$

Proof. By the assumption that $x_t := f(X, u^{(\psi_t)}(X)) \in \mathcal{M}_t$, this vector field lies on the manifold of the diffusion model at time t and therefore by Proposition 1, there exists a scaling constant $\bar{\alpha}_t$ and an asymptotically stable vector field $x \in \mathcal{M}$ such that $x_t = \sqrt{\bar{\alpha}_t}x + \sqrt{1 - \bar{\alpha}_t}\varepsilon$. As the diffusion model is given by the DDIM algorithm as stated in (4) and (5), this states that there exists an asymptotically stable vector field $x_{t-1} \in \mathcal{M}_{t-1}$, concentrated on \mathcal{M}_{t-1} with high probability, such that

$$x_{t-1} = \sqrt{\bar{\alpha}_{t-1}}x_{0|t} + \sqrt{1 - \bar{\alpha}_{t-1}} \frac{x_t - \sqrt{\bar{\alpha}_t}x_{0|t}}{\sqrt{1 - \bar{\alpha}_t}}. \quad (14)$$

Substituting $x_t = f(X, u^{(\psi_t)}(X))$ and $x_{t-1} = f(X, u(X))$ gives

$$\begin{aligned} f(X, u(X)) &= \frac{\sqrt{1 - \bar{\alpha}_{t-1}}}{\sqrt{1 - \bar{\alpha}_t}} f(X, u^{(\psi_t)}(X)) \\ &\quad + \left(\sqrt{\bar{\alpha}_{t-1}} - \frac{\sqrt{1 - \bar{\alpha}_{t-1}}\sqrt{\bar{\alpha}_t}}{\sqrt{1 - \bar{\alpha}_t}} \right) x_{0|t} \end{aligned} \quad (15)$$

By a first-order approximation,

$$f(X, u(X)) = f(X, u^{(\psi_t)}(X)) - \nabla_\psi f(X, u^{(\psi_t)}(X))h \quad (16)$$

for some vector $h \in \mathbb{R}^n$. Set

$$-h = [\nabla_{\psi} f(X, u^{(\psi_t)}(X)) \nabla_{\psi} f(X, u^{(\psi_t)}(X))^T]^{-1} \quad (17)$$

$$\nabla_{\psi} f(X, u^{(\psi_t)}(X))^T \quad (18)$$

$$\begin{aligned} & \left[\left(1 - \frac{\sqrt{1 - \bar{\alpha}_{t-1}}}{\sqrt{1 - \bar{\alpha}_t}}\right) f(X, u^{(\psi_t)}(X)) \right. \\ & \left. + \left(\sqrt{\bar{\alpha}_{t-1}} - \frac{\sqrt{1 - \bar{\alpha}_{t-1}} \sqrt{\bar{\alpha}_t}}{\sqrt{1 - \bar{\alpha}_t}} \right) x_{0|t} \right]. \end{aligned} \quad (19)$$

Since $f(X, u^{(\psi_t)}(X)) [\nabla_{\psi} f(X, u^{(\psi_t)}(X)) \nabla_{\psi} f(X, u^{(\psi_t)}(X))^T]^{-1}$ is the matrix representation of the projection onto the column space of $f(X, u^{(\psi_t)}(X))$ this is the projection onto the tangent space of \mathcal{W} at the point $f(X, u^{(\psi_t)}(X))$. Therefore, as r_t is assumed to lie on the tangent space, the projection matrix

$$\begin{aligned} P &= \nabla_{\psi} f(X, u^{(\psi_t)}(X)) \\ & [\nabla_{\psi} f(X, u^{(\psi_t)}(X)) \nabla_{\psi} f(X, u^{(\psi_t)}(X))^T]^{-1} \\ & \nabla_{\psi} f(X, u^{(\psi_t)}(X))^T \end{aligned} \quad (20)$$

acts as the identity matrix on r_t . It follows that

$$\nabla_{\psi} f(X, u^{(\psi_t)}(X)) h = \left(1 - \frac{\sqrt{1 - \bar{\alpha}_{t-1}}}{\sqrt{1 - \bar{\alpha}_t}}\right) r_t \quad (21)$$

and (15) is satisfied for this choice. Therefore, substituting and simplifying gives

$$f(X, u(X)) = f(X, u^{(\psi_t)}(X)) - c_t r_t \quad (22)$$

□

Remark 1. Note that since r_t lies on the tangent plane we see that x_{t-1} lies in $T_{x_t} \mathcal{W} \cap \mathcal{M}_{t-1}$. Therefore, by Assumption 3 there exists parameters ψ_{t-1} such that $f(x, u^{\psi_{t-1}}(X))$ lies in an ε neighborhood of $f(X, u(X))$ obtained from Proposition 2. Intuitively, this shows that updating parameters in (9) computes the difference in parameters corresponding to a linear shift from $\mathcal{M}_t \cap \mathcal{W}$ to \mathcal{M}_{t-1} .

Theorem 2. Let \mathcal{M} be the set of asymptotically stable vector fields projected onto a grid. Let c_t be defined as in equation (12). Assume the diffusion model ε_{θ} is optimal. With the update rule (8)–(10) with $L_t(x_t; \psi_t)$ defined as in equation (11), we can obtain an $x_{t-1} \sim N(x_{t-1}; f(X, u^{\psi_{t-1}}(X)), \sigma_t^2 I)$ whose marginal distribution

$$p(x_{t-1}) = \int \mathcal{N}(x_{t-1}; f(X, u^{\psi_{t-1}}(X)), \sigma_t^2 I) p(x_t | x) p(x) dx \quad (23)$$

is probabilistically concentrated on $\mathcal{M}_{t-1} \cap \mathcal{W}$

Proof. We first acknowledge that there is a one-to-one correspondence between x_t and ψ_t for all $t = 0, \dots, T$, thus showing that $p(x_t) = p(\psi_t)$. We also acknowledge that $x_t \in \mathcal{W}$ for all t since the update rule $x_t = f(X, u^{(\psi_t)}(X))$ forces this property on x_t .

We then prove that for all t , there exists an $x \in \mathcal{M} \cap \mathcal{W}$ such that the x_t generated from equations (8)–(10) can also be generated by the forward process of the diffusion model constrained to the \mathcal{W} manifold. In other words, $\psi_t = \sqrt{\bar{\alpha}_t} \psi + \sqrt{1 - \bar{\alpha}_t} \varepsilon$ for some $\psi \in \mathcal{W}$, $\varepsilon \in \mathcal{N}(0, I)$, where $x = f(X, u^{(\psi)}(X))$. We prove this using induction.

For the base case, let $t = T$. Since the parameters ψ_T are drawn from a Gaussian prior, $\psi_T = \sqrt{\bar{\alpha}_T} \psi + \sqrt{1 - \bar{\alpha}_T} \varepsilon = \varepsilon$ for any $\psi \in \mathcal{M}$ since $\bar{\alpha}_T = 0$.

The proof for $t \leq T$, has been done in Proposition 2 and Remark 1. We are guaranteed the existence of parameters ϕ_t for any $t \leq T$ such that $f(X, u^{\phi_t}(X))$ lies in an ε neighborhood of some $x_t \in \mathcal{M}_t$. This completes the proof by induction.

□

Proposition 2 and Theorem 2 allow us to propose a methodology for updating the parameters ψ to design a stabilizing controller.

3.3 Algorithm

We provided our proposed algorithm in Algorithm 1 below. This algorithm returns the stable vector field x_0 and the stabilizing controller ψ_0 . This algorithm will be tested in four baselines to derive stabilizing controllers.

Algorithm 1 Manifold Guided Lyapunov Control (MGLC)

Require: $f(\cdot, u(\cdot))$, Pretrained diffusion model $\varepsilon_\theta(\cdot, \cdot)$

- 1: Initialize $\psi_T \sim \mathcal{N}(0, I)$
 - 2: **for** $t=T, \dots, 1$ **do**
 - 3: Set $x_t = f(X, u(\psi_t))$
 - 4: Estimate $x_{0|t}$ using (4)
 - 5: Compute $L_t(x_t; \psi_t)$ and c_t using (11) and (12) respectively
 - 6: Update $\psi_{t-1} = \psi_t - c_t \nabla_{\psi_t} L_t(x_t; \psi_t)$
 - 7: **end for**
 - 8: **return** $x_0 = f(X, u(\psi_0))$ and ψ_0
-

4 IMPLEMENTATION

In this section, we explain the dataset and the unconditional diffusion model we trained for this paper. To ensure the reproducibility of our results, we provided our source code in the following GitHub repository: <https://github.com/amartyamukherjee/minim> diffusion

An overview of our model is posted in Figure 2. We train the diffusion model to output images with three channels. The first two channels correspond to $f_1(X)$ and $f_2(X)$, the two scalar elements of an asymptotically stable vector field. The third channel corresponds to $V(X)$, the Lyapunov function that verifies their stability.

4.1 Dataset Generation

For our diffusion model, we made a dataset of 2,000 pairs of (f, V) in 2D space. The first 1,000 pairs we made have the following format:

$$f(x) = Ax + \beta_f^T \tanh(W_f x) \quad (24)$$

$$V(x) = \beta_V \cdot \tanh(W_V x + b_V), \quad (25)$$

where A is a 2×2 Hurwitz matrix. Each entry of $\beta_f \in \mathbb{R}^{2 \times 20}$ is sampled from $\mathcal{N}(0, 0.2^2)$, and each entry of $W_f \in \mathbb{R}^{20 \times 2}$ is sampled from $\mathcal{N}(0, 1)$. $V(x)$ is a neural Lyapunov function trained identically to the method in [23] to minimize the loss $(\dot{V} + \|x\|^2)$.

For the next 1,000 pairs, we focus on second-order systems. We use the fact stated below.

Fact 1 (Example 4.8 of [11]). *Consider the system*

$$\dot{x}_1 = x_2 \quad (26)$$

$$\dot{x}_2 = -h_1(x_1) - h_2(x_2), \quad (27)$$

where $h_1(\cdot)$ and $h_2(\cdot)$ are locally Lipschitz and satisfy $h_i(0) = 0, y h_i(y) > 0, \forall y \in (-a, a) \setminus \{0\}$. This system is asymptotically stable with the following Lyapunov function candidate

$$V(x_1, x_2) = \int_0^{x_1} h_1(y) dy + \frac{1}{2} x_2^2, \quad (28)$$

satisfying

$$\dot{V} = -x_2 h_2(x_2). \quad (29)$$

Using this fact, we can design the following system:

$$\dot{x}_1 = x_2 \quad (30)$$

$$\dot{x}_2 = -c_1 x_1 - c_2 \tanh x_1 - c_3 x_2 - c_4 \tanh x_2 \quad (31)$$

$$V(x_1, x_2) = \frac{c_1}{2} x_1^2 + c_2 \log(\cosh x_1) + \frac{1}{2} x_2^2, \quad (32)$$

where c_i is drawn from the uniform distribution $U(0, 5)$. This ensures that $\dot{V} = -x_2(c_3 x_2 + c_4 \tanh x_2) \leq -c_3 x_2^2$.




Figure 2: An overview of our model in the 2D control setting. We train the generator to output images with three channels. The first two channels are the scalar elements of an asymptotically stable vector field, $f_1(x,y)$ and $f_2(x,y)$. The third channel is a Lyapunov function $V(x,y)$ that verifies the stability of the vector field.

4.2 Diffusion model

We trained a diffusion model by modifying the GitHub repository by [17] that is based on DDIM. This model involves a forward (or "diffusion process") that is a Markov chain, which gradually adds Gaussian noise to the data given by a cosine scheduler. Upon training our diffusion model, we needed to normalize our vector fields and Lyapunov function so that each pixel has values in $[-1, 1]$. This means that our model is trained with Lyapunov functions that satisfy $\dot{V} = c\|x\|^2$ for the first 1,000 data points and $\dot{V} < cx_2^2$ for the second 1,000 data points, where c is a strictly positive constant.

We trained our diffusion model in parallel using two NVIDIA 3090Ti GPUs for our dataset of 2,000 samples and the training time was approximately 8 hours.

4.3 Controller

The format of the controller is

$$u^{(\psi)}(x) = C \sum_{i=1}^n \tanh(\psi_i x_i), \quad (33)$$

where $\psi_i, i = 1, \dots, n$ are trainable parameters. At every step of the reverse diffusion process, we update the parameters ψ using the loss function defined in equation (11) and the step size in equation (12).

5 NUMERICAL RESULTS

We finally used this pre-trained model to derive stabilizing controllers for four unseen nonlinear systems. We notice that the derivation of a controller took only 16 seconds.

5.1 Inverted Pendulum

The inverted pendulum is governed by the following dynamics:

$$\ddot{\theta} = \frac{mgl \sin \theta + u - 0.1 \dot{\theta}}{ml^2}, \quad (34)$$

where our constants are $g = 9.81, m = 0.15, l = 0.5$. Using our MGLC algorithm, our control function is

$$u(x_1, x_2) = 20 \tanh(-4.16928\theta) + 20 \tanh(-3.14848\dot{\theta}). \quad (35)$$

To verify the stability of this system, we solved it numerically with 100 different initial conditions using RK45 and plotted the trajectories in Figure 4. The plot shows that all the trajectories converged to the origin, thus showing that our control function is stabilizing.




Figure 3: Visualization of MGLC for the inverted pendulum environment




Figure 4: Trajectories of the controlled inverted pendulum system with 100 different randomly sampled initial conditions

We have also posted a visualization of the convergence of our system for 5 time steps in Figure 3. Since $f_{1,t}(x) = x_2$, we decided to omit that channel in this visualization. We have also plotted our estimate of the second channel of $x_{0|t}$, which we will call $f_{2,0|t}$. That channel shows a convergence of our system to a stable but noisy vector field, which is used in our loss function for updating ψ . We have also shown the evolution of our candidate Lyapunov function in the third channel.

5.2 Damped Duffing Oscillator

The damped Duffing oscillator is governed by the following dynamics:

$$\dot{x}_1 = x_2 \quad (36)$$

$$\dot{x}_2 = -0.5x_2 - x_1(4x_1^2 - 1) + 0.5u. \quad (37)$$

The uncontrolled system ($u = 0$) has two stable equilibrium points, $(-0.5, 0)$ and $(0.5, 0)$, and an unstable equilibrium point in the origin $(0, 0)$. Using our MGLC algorithm, our control function is

$$u(x_1, x_2) = 20 \tanh(-3.89859x_1) + 20 \tanh(-4.46941x_2). \quad (38)$$

The numerical solutions of our system, plotted in Figure 5, show that our control function is stabilizing.




Figure 5: Trajectories of the controlled damped Duffing oscillator system with 100 different randomly sampled initial conditions

5.3 Van Der Pol Oscillator

The controlled Van Der Pol oscillator system is governed by the following dynamics:

$$\dot{x}_1 = 2x_2 \quad (39)$$

$$\dot{x}_2 = -0.8x_1 + 2x_2 - 10x_1^2x_2 + u. \quad (40)$$

The uncontrolled system ($u = 0$) shows a limit cycle in a neighborhood near the origin, and an unstable equilibrium point in the origin $(0, 0)$. Using our MGLC algorithm, our control function is

$$u(x_1, x_2) = 20 \tanh(-5.05384x_1) + 20 \tanh(-3.25052x_2). \quad (41)$$

The numerical solutions of our system, plotted in Figure 6, show that our control function is stabilizing.




Figure 6: Trajectories of the controlled Van Der Pol oscillator system with 100 different randomly sampled initial conditions

5.4 Noisy Inverted Pendulum

In this experiment, we consider the pendulum system in equation (34) with uncertainties in the parameters m, l and the control u . At each time step, we replace m, l in the system with $m + z_1, l + z_2$, and the control is given by $(1 + z_3)u^{(\psi)}(x)$, where z_i are sampled from the uniform distribution $U(-0.05, 0.05)$ for $i = 1, 2, 3$.

$$u(x_1, x_2) = 20 \tanh(-4.01703x_1) + 20 \tanh(-3.63485x_2). \quad (42)$$

To verify the stability of this system, we solved it numerically with 100 different initial conditions using the Euler-Maruyama scheme and plotted the trajectories in Figure 7. The plot shows that all the trajectories converged to the origin, thus showing that our control function is stabilizing.

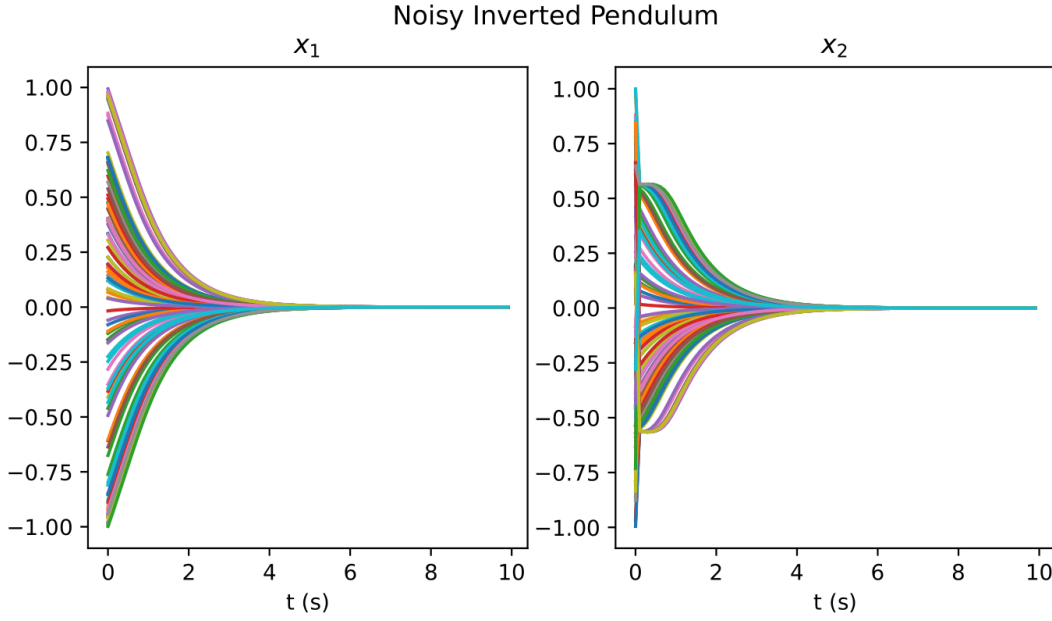


Figure 7: Trajectories of the controlled inverted pendulum system designed with uncertainties in parameters and control inputs on 100 different randomly sampled initial conditions

6 CONCLUSIONS AND FUTURE WORKS

6.1 Conclusions

This paper introduces a new approach to the control of nonlinear systems by identifying the closest asymptotically stable vector field relative to a predetermined manifold. We employ diffusion models to aid us in estimating this stable vector field using Tweedie’s estimate. We finally use this estimate to update the parameters of our control function. This method helps us achieve fast zero-shot control for control problems that are unseen by the diffusion model. This work, as a result, presents a novel application of diffusion models for transfer learning in control problems. Furthermore, the design of a stabilizing control only needs evaluations of the vector field $f(x, u)$ at grid points X , thus showing that this can easily be implemented as a data-driven approach.

6.2 Future Works

One key area of further exploration is the extension of this work to systems of high dimensions. Since the diffusion model outputs images in a pixel format, we were able to only work with 2D systems. The extension to high dimensions may be possible using models such as latent diffusion [16]. The work of [7] has provided improved convergence guarantees using perfect encoders as done in latent diffusion. We expect that, equivalently, improved convergence guarantees can be expected using perfect encoders of high-dimensional systems.

7 ACKNOWLEDGMENTS

We sincerely acknowledge the support of Professor Yaoliang Yu from the Department of Computer Science, University of Waterloo for providing us with GPUs to train our model and for his valuable feedback on the paper.

References

- [1] Daniele Ahmed, Andrea Peruffo, and Alessandro Abate. Automated and sound synthesis of lyapunov functions with smt solvers. In *Proc. of TACAS*, pages 97–114. Springer, 2020.
- [2] Rucha Apte, Sheel Nidhan, Rishikesh Ranade, and Jay Pathak. Diffusion model based data generation for partial differential equations. *ICML workshop on Synergy of Scientific and Machine Learning Modeling*, 2023.

- [3] Ya-Chien Chang, Nima Roohi, and Sicun Gao. Neural lyapunov control. *Advances in neural information processing systems*, 32, 2019.
- [4] Hyungjin Chung, Jeongsol Kim, Michael Thompson Mccann, Marc Louis Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. In *International Conference on Learning Representations*, 2023.
- [5] Hyungjin Chung, Byeongsu Sim, Dohoon Ryu, and Jong Chul Ye. Improving diffusion models for inverse problems using manifold constraints. *Advances in Neural Information Processing Systems*, 35:25683–25696, 2022.
- [6] Bradley Efron. Tweedie’s formula and selection bias. *Journal of the American Statistical Association*, 106(496):1602–1614, 2011.
- [7] Yutong He, Naoki Murata, Chieh-Hsin Lai, Yuhta Takida, Toshimitsu Uesaka, Dongjun Kim, Wei-Hsiang Liao, Yuki Mitsufuji, J Zico Kolter, Ruslan Salakhutdinov, et al. Manifold preserving guided diffusion. *International Conference on Learning Representations*, 12, 2024.
- [8] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [9] Emiel Hooeboom, Victor Garcia Satorras, Clément Vignac, and Max Welling. Equivariant diffusion for molecule generation in 3d. In *International conference on machine learning*, pages 8867–8887. PMLR, 2022.
- [10] Bahjat Kavar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. In *Advances in Neural Information Processing Systems*, 2022.
- [11] Hassan K. Khalil. *Nonlinear Systems, Third Edition*. Pearson Education, 2002.
- [12] Jun Liu, Yiming Meng, Maxwell Fitzsimmons, and Ruikun Zhou. Physics-informed neural network lyapunov functions: Pde characterization, learning, and verification. *arXiv preprint arXiv:2312.09131*, 2023.
- [13] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in Neural Information Processing Systems*, 35:5775–5787, 2022.
- [14] Naoki Murata, Koichi Saito, Chieh-Hsin Lai, Yuhta Takida, Toshimitsu Uesaka, Yuki Mitsufuji, and Stefano Ermon. GibbsDDRM: A partially collapsed gibbs sampler for solving blind inverse problems with denoising diffusion restoration. In *International Conference on Machine Learning*, 2023.
- [15] Oded Ovadia, Eli Turkel, Adar Kahana, and George Em Karniadakis. Ditto: Diffusion-inspired temporal transformer operator. *arXiv preprint arXiv:2307.09072*, 2023.
- [16] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022.
- [17] Vikash Sehwal. minimal-diffusion, <https://github.com/VSehwal/minimal-diffusion>, 2022.
- [18] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *International Conference on Learning Representations*, 2021.
- [19] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.
- [20] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.
- [21] Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. *ACM Computing Surveys*, 56(4):1–39, 2023.

- [22] Ye Yuan, Jiaming Song, Umar Iqbal, Arash Vahdat, and Jan Kautz. Physdiff: Physics-guided human motion diffusion model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16010–16021, 2023.
- [23] Ruikun Zhou, Thanin Quartz, Hans De Sterck, and Jun Liu. Neural lyapunov control of unknown nonlinear systems with stability guarantees. *Advances in Neural Information Processing Systems*, 35:29113–29125, 2022.