

Globally Stable Neural Imitation Policies

Amin Abyaneh¹, Mariana Sosa Guzmán², and Hsiu-Chin Lin³

Abstract—Imitation learning mitigates the resource-intensive nature of learning policies from scratch by mimicking expert behavior. While existing methods can accurately replicate expert demonstrations, they often exhibit unpredictability in unexplored regions of the state space, thereby raising major safety concerns when facing perturbations. We propose SNDS, an imitation learning approach aimed at efficient training of scalable neural policies while formally ensuring global stability. SNDS leverages a neural architecture that enables the joint training of the policy and its associated Lyapunov candidate to ensure global stability throughout the learning process. We validate our approach through extensive simulations and deploy the trained policies on a real-world manipulator arm. The results confirm SNDS’s ability to address instability, accuracy, and computational intensity challenges highlighted in the literature, positioning it as a promising solution for scalable and stable policy learning in complex environments.

I. INTRODUCTION

Imitation learning (IL) is a pivotal notion attempting to overcome the diverse safety and complexity challenges of policy learning by emulating expert behavior [1], [2]. IL facilitates the training of intricate motion policies without resorting to an exhaustive search in the robot’s state space [3], [4]. Nonetheless, only a handful of IL methods offer formal guarantees pertaining to the stability of the resulting policies. Global stability assumes a critical role when deploying the policy in stochastic environments prone to external perturbations. Notably, it assures that the policy can recover effectively and *predictably* to a predetermined target, even in uncharted regions of the state space that lie beyond the scope of expert demonstrations.

Prior research in safe IL is focused on autonomous and time-invariant dynamical systems to model and train motion policies [4]. These policies are optimized across a set of expert trajectories, yielding a proper action (velocity) contingent on the current state (position). The optimization process ensures global stability through established theoretical frameworks, such as Lyapunov or contraction theories [5]–[11]. Nevertheless, reproducing complex trajectories in high-dimensional state spaces with the previous methods can be impractical due to computational complexity, non-convex optimization, and sample inefficiency. To address

¹Amin Abyaneh is with the Department of Electrical and Computer Engineering, McGill University, Montreal, Canada, amin.abbyaneh@mail.mcgill.ca

²Mariana Sosa Guzmán is with the Department of Electrical and Computer Engineering, Universidad Veracruzana, Veracruz, Mexico.,

³Hsiu-Chin Lin is with the School of Computer Science and the Department of Electrical and Computer Engineering, McGill University, Montreal, Canada, hsiu-chin.lin@cs.mcgill.ca

This work is sponsored by NSERC Discovery Grant and Mitacs Globalink Research Internship.

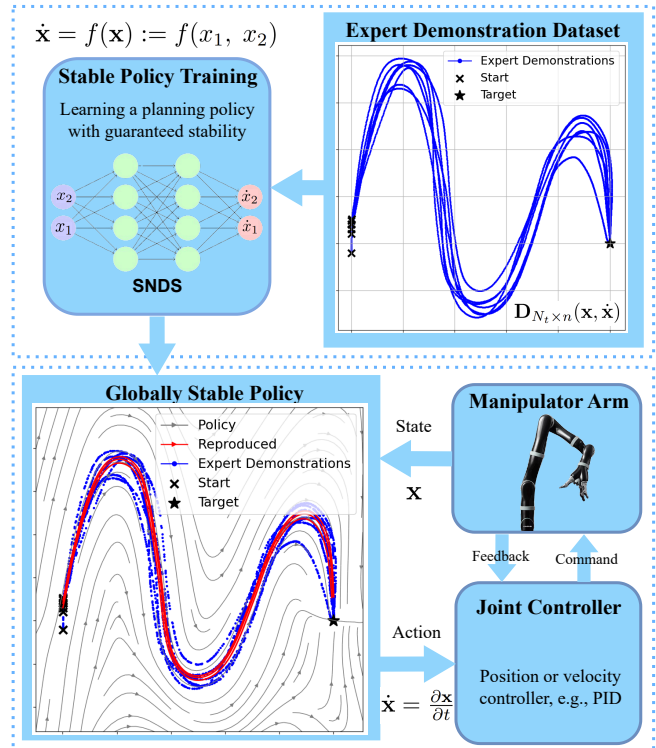


Fig. 1. Overview of the proposed stable neural policy learning method. Policy learning (top) optimizes a Lyapunov-stable neural policy over the expert demonstration data. The optimized policy is then deployed (bottom) to plan globally stable trajectories resistant to unpredictable perturbations.

computational cost, previous research [7], [9], [11] tends to restrict the class of functions of the contraction metric or the Lyapunov candidate, which in turn limits the imitation performance of the optimized policy. Most notably, a common limitation of these methods appears in the absence of neural policy representation, which are recognized for their scalability, efficient gradient-based optimization, and domain transfer capabilities [12].

Recently developed stable neural policies [13], [14] mainly employ diffeomorphism and invertible mappings to transform a simple stable policy into a highly nonlinear one. Yet, the training stage requires multiple demonstrations, leads to policies with quasi-stability, and hinders incremental learning for slightly different expert trajectories. Neural IL methods, on the other hand, offer flexible and precise policies [15]–[18], but lack the required safety and reliability outside the region covered by expert demonstrations and are often data hungry. Similar limitations persist in inverse reinforcement learning [19], [20], where computational complexity is fueled

by having reinforcement learning in an inner loop [15].

To tackle these challenges, we represent policies with **Stable Neural Dynamical Systems (SNDS)**, where we jointly train a neural dynamical system policy, alongside a secondary, convex-by-design network that guarantees global stability. Subsequently, a joint optimization process trains the neural policy by minimizing a novel differentiable loss aimed at strengthening the alignment between policy rollouts and expert demonstrations. SNDS, therefore, benefits from an expressive and stable neural representation, allowing for safe and scalable approximation of the underlying dynamical system. We present an overview of our framework in Figure 1, and outline our key **contributions** below.

- Designing SNDS—a *stable-by-design* neural representation for nonlinear dynamical systems as motion policies based on expert demonstrations.
- Providing formal stability analysis founded on Lyapunov theory and convex neural networks.
- Formulating a differentiable trajectory alignment loss function, inspired by forward Euler’s method.
- Empirical evaluation of SNDS’s effectiveness in higher state space dimensions for complex trajectories, both in simulation and real-world scenarios.

II. BACKGROUND

A. Preliminaries

Dynamical system. A dynamical system (DS) describes the evolution of a state, \mathbf{x} , in the ambient space $\mathcal{X} \subset \mathbb{R}^n$, over time [21]. We consider an autonomous, time-invariant DS modeled with a first-order ordinary differential equation, $\dot{\mathbf{x}} = f(\mathbf{x})$, where the system yields the time derivative for each \mathbf{x} , without any control input and independent of time.

Lyapunov stability theory. $f(\mathbf{x})$ is globally asymptotically stable (GAS) at an equilibrium, \mathbf{x}^* , if for any initial state, the system approaches \mathbf{x}^* as time progresses towards infinity. The Lyapunov stability theory is a widely used tool to analyze the GAS property of dynamical systems. According to the theory, a system exhibits GAS if there exists a positive-definite function $\psi(\mathbf{x}) : \mathcal{X} \rightarrow \mathbb{R}$, referred to as the Lyapunov potential function (LPF), satisfying $\dot{\psi}(\mathbf{x}) < 0$ for all $\mathbf{x} \neq \mathbf{x}^*$, and $\psi(\mathbf{x}^*) = 0$ at the equilibrium.

Input convex neural networks (ICNN). ICNN’s special architecture enables universal approximation of convex functions while maintaining convexity during training [22], [23]. An L -layer ICNN, $\hat{\psi}$, is formulated as follows for $l \in [L]$:

$$\hat{\psi}(\mathbf{x}) = z_L, \quad z_{l+1} = \sigma_l(U_l x + W_l z_l + b_l),$$

where z_l denotes the layer output, and W_l and b_l are real-value weights. In contrast to feedforward networks, ICNN design exploits U_l , that are positive weights connecting the input directly to each layer, and σ_l can be any convex, non-decreasing activation function.

B. Problem statement

Consider a policy that functions within a state space denoted by $\mathcal{X} \subset \mathbb{R}^n$. The state space may correspond to a robot’s task

(T) or configuration (C) space. The policy outputs an action in $\mathcal{A} \subset \mathbb{R}^n$, such as velocity or torque, which determines the change in the state over time. We denote the state variable by $\mathbf{x} \triangleq [x_1; x_2; \dots; x_n]^T \in \mathcal{X}$, and assume that the yielded action indicates the time derivative of the current state, denoted by $\dot{\mathbf{x}} = \frac{\partial \mathbf{x}}{\partial t} \in \mathcal{A}$. In this context, our **primary objective** is to learn a globally stable imitation policy to map \mathbf{x} to $\dot{\mathbf{x}}$, provided a dataset of expert’s state-action pairs.

Let $N_d \in \mathbb{N}$ be the number of trajectories demonstrated by the expert. Each trajectory contains $N_s \in \mathbb{N}$ state-actions. The dataset of expert trajectories stacks all state-action pairs,

$$\mathbf{D} \triangleq \left\{ (\mathbf{x}^d[s], \dot{\mathbf{x}}^d[s]) \mid d \in [N_d], s \in [N_s] \right\}, \quad (1)$$

where $(\mathbf{x}^d[s], \dot{\mathbf{x}}^d[s])$ is the dataset entry corresponding to the s -th sample of the d -th trajectory. The dataset \mathbf{D} holds $N_t = N_d N_s$ samples. We assume that the trajectories share a common target endpoint, $\mathbf{x}^* \in \mathcal{X}$, and have zero velocity at the target, i.e., $\mathbf{x}^d[N_s] = \mathbf{x}^*$ and $\dot{\mathbf{x}}^d[N_s] = \mathbf{0}$, $\forall d \in [N_d]$.

The mapping between states and actions can be modeled with a time-invariant autonomous DS [5], denoted by:

$$\dot{\mathbf{x}} = f(\mathbf{x}) + \varepsilon = \phi(\mathbf{x}), \quad f, \phi : \mathcal{X} \rightarrow \mathcal{A}. \quad (2)$$

In Equation (2), the function f corresponds to an ordinary differential equation describing the true underlying DS. The term $\varepsilon \in \mathbb{R}^n$ accounts for the effect of measurement noise present in the expert’s demonstrations, which is incorporated into the estimated DS. Our subsequent objective is to acquire a noise-free approximation of $f(\mathbf{x})$, denoted as $\phi(\mathbf{x})$. This estimated function, $\phi(\mathbf{x})$, can be perceived as a *policy* that maps states to actions in real-time, thereby guiding the robot with the right set of actions to imitate the demonstrated trajectories within the state space.

Alternatively, a trajectory can be generated through forward Euler’s method. Imagine a robot situated at $\mathbf{x}[0]$, the policy generates an action $\dot{\mathbf{x}}[0] = \phi(\mathbf{x}[0])$, and the next state is subsequently calculated through $\mathbf{x}[1] = \mathbf{x}[0] + \Delta t \phi(\mathbf{x}[0])$, where Δt determines the granularity of the discretization.

III. METHODOLOGY

We model the policy, $\phi(\mathbf{x})$, and the corresponding LPF, $\psi(\mathbf{x})$, with two neural networks in Section III-A. This allows us to accurately imitate an expert’s behavior, while providing formal GAS analysis in Section III-B. Lastly, we introduce a differentiable loss function in Section III-C to improve both sample efficiency and imitation accuracy of SNDS.

A. Dynamical system policy formulation

Let $\hat{\phi}(\mathbf{x}) : \mathcal{X} \rightarrow \mathcal{A}$ be a standard and unrestricted feedforward neural network which models the unstable policy. Although the model can capture intricate expert demonstrations, training $\hat{\phi}(\mathbf{x})$ solely on \mathbf{D} results in an unstable policy which cannot predictably recover even from mild perturbations, as depicted in Figure 2. Hence, in this section, we enforce the trained policy to have the GAS property by satisfying the conditions of Lyapunov stability theory in Section II-A.

The first step is to define a positive-definite and differentiable LPF. As mentioned in Section II-A, ICNNs can

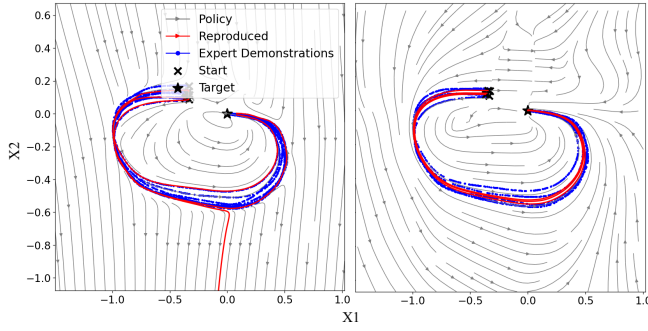


Fig. 2. An example of unstable (left) vs. a stable (right) policies optimized on expert’s data from the handwriting dataset [24]. While both policy rollouts can reproduce the expert motion, an unstable policy cannot recover from perturbation that push the robot to unknown state space regions.

approximate any convex function. To ensure that the LPF is positive definite, we employ a well-studied technique [17], [18] to define it as,

$$\psi(\mathbf{x}) = [\hat{\psi}(\mathbf{x}) - \hat{\psi}(\mathbf{x}^*)] + \delta \|\mathbf{x} - \mathbf{x}^*\|_2^2, \quad (3)$$

where $\psi(\mathbf{x}) : \mathcal{X} \rightarrow \mathbb{R}$ is still a convex function. The first term ensures $\psi(\mathbf{x}^*) = 0$, while the addition of the state’s l_2 -norm with a negligible $\delta > 0$, along with convexity of ICNN guarantee that $\psi(\mathbf{x}) > 0, \forall \mathbf{x} \neq \mathbf{x}^*, \mathbf{x} \in \mathcal{X}$. Hence, the LPF with this architecture, $\psi(\mathbf{x})$, satisfies the positive-definite condition required by Lyapunov stability theory.

Next, to satisfy the negative derivative condition of Lyapunov theory, i.e., $\frac{d\psi(\mathbf{x})}{dt} = \nabla\psi(\mathbf{x})^T \hat{\phi}(\mathbf{x}) < 0$, we modify the projection expression in [23] to enforce GAS while training $\hat{\phi}(\mathbf{x})$. The policy $\phi(\mathbf{x})$ resulting from projecting $\hat{\phi}(\mathbf{x})$ onto the half space, $\nabla\psi(\mathbf{x})^T \hat{\phi}(\mathbf{x}) < 0$, is formulated as follows,

$$\begin{aligned} \phi(\mathbf{x}) &= \begin{cases} \hat{\phi}(\mathbf{x}) - \nabla\psi(\mathbf{x}) \frac{\nabla\psi(\mathbf{x})^T \hat{\phi}(\mathbf{x})}{\|\nabla\psi(\mathbf{x})\|_2} & \nabla\psi(\mathbf{x})^T \hat{\phi}(\mathbf{x}) \geq 0 \\ \hat{\phi}(\mathbf{x}) & \nabla\psi(\mathbf{x})^T \hat{\phi}(\mathbf{x}) < 0 \end{cases} \\ &= \hat{\phi}(\mathbf{x}) - \nabla\psi(\mathbf{x}) \frac{\text{ReLU}(\nabla\psi(\mathbf{x})^T \hat{\phi}(\mathbf{x}))}{\|\nabla\psi(\mathbf{x})\|_2}. \end{aligned} \quad (4)$$

In Equation (4), if the network output $\hat{\phi}(\mathbf{x})$ fails to meet the Lyapunov condition, i.e., $\nabla\psi(\mathbf{x})^T \hat{\phi}(\mathbf{x}) \geq 0$, the output is projected so that $\phi(\mathbf{x})$ always fulfills the condition. The process is simplified with a ReLU activation function. Note that the original formulation in [23] enforces exponential stability, which is too restrictive for our problem. Hence, we further relax the projection to guarantee GAS.

Figure 3 portrays an example of trained $\phi(\mathbf{x})$ and $\psi(\mathbf{x})$, and the resulting global stability. The differentiable and globally stable $\phi(\mathbf{x})$, can be trained using efficient gradient-based methods within its parameters, as long as $\hat{\phi}(\mathbf{x})$ and $\psi(\mathbf{x})$ are defined using automatic differentiation tools.

B. Global asymptotic stability guarantees

To establish that $\phi(\mathbf{x})$ defines a function with GAS property, we use proof by contradiction to show that regardless of initial state, $\mathbf{x}[0]$, every trajectory converges to the target, that is, $\lim_{k \rightarrow \infty} \mathbf{x}[k] = \mathbf{x}^*$.




Fig. 3. Joint training of globally stable neural policies (left) and the corresponding LPF (right) optimized on expert’s data from the handwriting dataset [24]. The contours for the LPF (green) illustrate both positive-definiteness and convexity of the trained function.

Proposition 1. *The dynamical system, $\phi(\mathbf{x})$, in Equation (4) is globally asymptotically stable with the Lyapunov function, $\psi(\mathbf{x})$, defined in Equation (3), and any two arbitrary networks, $\hat{\phi}(\mathbf{x})$ and $\hat{\psi}(\mathbf{x})$, with bounded real-valued weights.*

Proof: We leverage proof by contradiction for this proposition. Let $\{\mathbf{x}[0], \dots, \mathbf{x}[k]\}$ be a trajectory not converging to the target, meaning that $\lim_{k \rightarrow \infty} \mathbf{x}[k] \neq \mathbf{x}^*$. Since $\psi(\mathbf{x})$ is decreasing, as formulated in Equation (4), and $\psi(\mathbf{x}) \geq 0$, as defined in Equation (3), there must exist a value $\mu \in \mathbb{R}^+$ such that $\lim_{k \rightarrow \infty} \psi(\mathbf{x}[k]) = \mu$. We reason that $\mu \leq \psi(\mathbf{x}[i]) \leq \psi(\mathbf{x}[0])$ for every $0 \leq i \leq k$.

Now consider the set \mathbf{S} of all $\mathbf{x}[i], \forall 0 \leq i \leq k$. This set is *compact*. More precisely, for any open cover¹ of that set, there exists a finite subcover, which means that a finite number of the open sets from the cover are sufficient to cover the original set. Hence, $\psi(\mathbf{x})$ takes its supremum, $\sup_{\mathbf{S}} \psi(\mathbf{x}) = -\alpha$ and $\alpha \in \mathbb{R}^+$ over this set. This conclusion is justified based on the projection introduced in Equation (4). The core integration in the proof of Lyapunov theorem,

$$\psi(\mathbf{x}[T]) = \psi(\mathbf{x}[0]) + \int_0^T \dot{\psi}(\mathbf{x}[t]) dt \leq \psi(\mathbf{x}[0]) - \alpha T,$$

indicates a *contradiction* when, $T = \frac{\psi(\mathbf{x}[0])}{\alpha} + \beta$ and $\beta \in \mathbb{R}^+$:

$$\psi(\mathbf{x}[T]) \leq \psi(\mathbf{x}[0]) - \alpha \left(\frac{\psi(\mathbf{x}[0])}{\alpha} + \beta \right) = -\alpha\beta < 0,$$

The above equation contradicts the fact that $\psi(\mathbf{x}) > 0$. Hence, all trajectories must converge to the desired target, \mathbf{x}^* . \square

Remark 2. *Proposition 1 only requires $\psi(\mathbf{x})$ to be convex. The policy network, $\phi(\mathbf{x})$, remains fairly unrestricted to capture intricate expert demonstrations.*

Thus, the proposed SNDS formally establishes global stability using an efficient gradient-based search for the LPF within the class of convex functions.

C. SRVF training loss

When constructing an optimization loss, the literature often opts for the conventional Mean Squared Error (MSE) [3],

¹An open cover for a set is a collection of open sets whose union contains the original set.

[5], [13] to gauge the dissimilarity between the policy’s output and the actual velocity labels over each batch of data. Consequently, the MSE loss is solely focused on the precision of estimated velocity. While reproducing exact velocities is essential, even slight inaccuracies at one instance can lead to accumulated error in the reproduced trajectories. The accumulation happens as a result of using forward Euler’s method that relies on previous estimates of \mathbf{x} to yield the next ones. Therefore, a more effective loss function should promote the accuracy of trajectories generated through forward Euler’s method during policy rollouts.

Employing forward Euler’s method to generate a full trajectory makes the gradient-based optimization intractable as a result of repetitive integrations. To mitigate this effect, we propose a limited horizon combination of the Square-Root Velocity Functions (SRVF) [25] and traditional MSE to design the loss function. SRVF methods involve normalizing the curve’s velocity, with a focus on changes in curve shape. Assuming a trajectory $\mathbf{x}_k^d = \{\mathbf{x}^d[0], \dots, \mathbf{x}^d[k]\}$, we shape the features by the definition:

$$\text{SRVF}(\mathbf{x}_k^d) = \sqrt{\frac{\mathbf{x}_k^d[s + \Delta t] - \mathbf{x}_k^d[s]}{\|\mathbf{x}_k^d[s + \Delta t] - \mathbf{x}_k^d[s]\|}}, \quad \Delta t = 1. \quad (5)$$

Given the state-action pairs, $\mathbf{x}^d[s]$, $\dot{\mathbf{x}}^d[s] \sim \text{SRVF}(\mathbf{D})$, and the differentiable dynamics $\phi(\mathbf{x})$, we define the loss \mathcal{L} as,

$$\begin{aligned} \mathcal{L}(\mathbf{x}; \theta) = & \gamma_0 \mathbb{E} \left[(\phi(\mathbf{x}^d[s]) - \dot{\mathbf{x}}^d[s])^2 \right] + \\ & \sum_{i \in [N_w]} \gamma_i \mathbb{E} \left[(\mathbf{x}^d[s + i - 1] + \phi(\mathbf{x}^d[s + i - 1])\Delta t - \mathbf{x}^d[s + i])^2 \right], \end{aligned} \quad (6)$$

where γ_i are discount factors decreasing as the loss horizon expands. The intuition behind Equation (6) is rather straightforward: \mathcal{L} is a differentiable function encapsulating both the current label mismatch, and deviation from the generated trajectory for a fixed horizon of N_w consecutive samples. As long as the horizon is limited, stochastic optimizers, such as ADAM [26], can optimize the loss on expert’s training data.

IV. EXPERIMENTS

We utilize two sets of motion planning data in our experiments. Our primary dataset is sourced from the well-known LASA Handwriting dataset [24], which contains records of handwritten trajectories on a tablet. The second dataset consists of more complex motions gathered by [7] in the same way. Our real-world experiments are only based on the latter due to its complex nature.

A. Evaluation

We compare our approach against three existing baselines: Behavioral Cloning (BC) [27], [28], Linear Parameter-Varying Dynamical System (LPV-DS) [7], and Stable Dynamical System Learning Using Euclideanizing Flows (SDS-EF) [13]. Among these methods, only BC lacks formal global asymptotic stability guarantees. Nevertheless, comparing SNDS against BC enables us to gauge SNDS’s accuracy and global stability against an unrestricted neural policy.

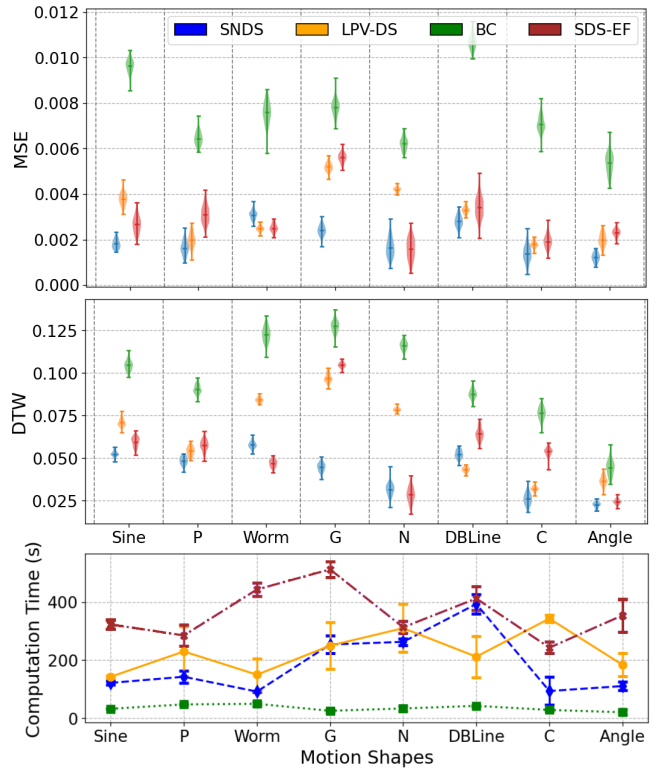


Fig. 4. Comparing the reproduction accuracy (top) and computational cost (bottom) of SNDS against baseline methods using the MSE and DTW metrics introduced in Equation (7) and Equation (8), respectively. The accuracy of policy actions and the learning time against other *stable* methods remain comparably lower for SNDS across most designated trajectories.

To assess the effectiveness of SNDS in comparison with the named baselines, we carry out the learning process for each method on designated handwriting dataset motions. For each motion, we randomly split the demonstrated dataset in the dataset into train (0.8) and test (0.2) sets. The policy learning stage is carried out on the training data, and the optimal policy is used to generate policy rollouts, \mathbf{x}^g , using the Euler’s method. Policies are subsequently evaluated by calculating Mean-Squared Error (MSE),

$$\text{MSE}(\dot{\mathbf{x}}^g, \dot{\mathbf{x}}^d) = \frac{\sum_{d=1}^{N_d^{\text{test}}} \sum_{s=1}^{N_s} (\phi(\mathbf{x}^d[s]) - \dot{\mathbf{x}}^d[s])^2}{2N_d^{\text{test}} N_s}, \quad (7)$$

and Dynamic Time Warping (DTW),

$$\text{DTW}(\mathbf{x}^g, \mathbf{x}^d) = \min_{p \in \mathcal{P}(\mathbf{x}^g, \mathbf{x}^d)} \left(\sum_{(i,j) \in \pi} \text{Euc}(\mathbf{x}^g[i], \mathbf{x}^d[j])^q \right)^{\frac{1}{q}}, \quad (8)$$

for handwriting dataset and experiments in higher state space dimensions to better capture discrepancies between policy rollouts and expert demonstrations. Within Equation (8), p represents an alignment path, \mathcal{P} is the set of all admissible paths, Euc is Euclidean distance [29], and $q = 2$. Notice that to calculate DTW, we need to generate an entire trajectory, while with MSE, only the generated actions are compared against true labels. We conduct the training and assessment




Fig. 5. Policy rollouts for SNDS and other baselines. Notice the highlighted inaccuracies or instabilities (yellow) for other methods. The acquired policies are optimized using the N-shaped data of the handwriting dataset.

procedures repeatedly with 10 random seeds, and present the mean and standard deviation of the results.

B. Handwriting dataset policies

We apply SNDS and selected baselines to each set of demonstrations as explained in Section IV-A, and compare the reproduction accuracy and computation time² across various motions in the handwriting dataset. Figure 4 presents a numerical comparison of both reproduction accuracy and computation time between SNDS and the selected baselines. The results indicate an acceptable level of reproduction accuracy while providing GAS certificates. In rare cases, SNDS shows a higher error compared to baselines, but outperforms others for the majority of elected motions, and maintains a reasonable computation time.

To closely examine the GAS property of each method, we plot the acquired policies using streamlines and generated two simulated rollouts with Euler’s method in Figure 5. It becomes apparent that SNDS, and LPV-DS provide predictability through guaranteeing GAS, while other approaches, such as BC and SDS-EF fail to render similar certificates. On a general note, even though we observe unstable policies trained with SDS-EF, instabilities occur only for a subset of motions in the handwriting dataset, and tend to intensify in regions further away from expert’s demonstrations.

To further showcase the sample inefficiency associated with Gaussian mixture model and diffeomorphism-based methods, such as LPV-DS and SDS-EF, respectively, we repeat the experiments with the training set reduced to only a *single demonstration*. The learned policies in this scenario

²Note that we restrict our computational resources to a single Core-i7 Gen8 CPU, despite that SNDS and SDS-EF can leverage parallel computing.

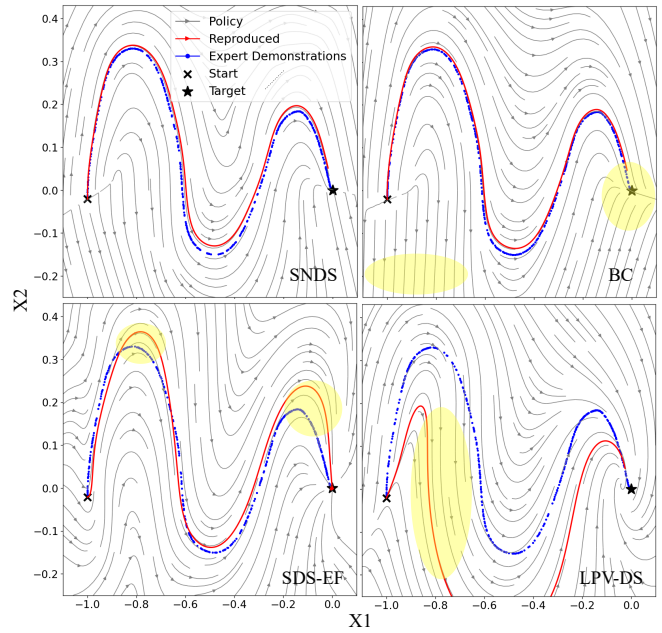


Fig. 6. Policy rollouts when training on a single Sine-shaped demonstration of the handwriting data. Inaccuracies or instabilities are highlighted. SNDS performs better owing to the expressive architecture and customized loss.

are illustrated in Figure 6, indicating the sample efficiency of SNDS when compared to LPV-DS and SDS-EF in the presence of limited demonstration samples.

C. SE(3) policy training

To illustrate the applicability of our method in high-dimensional state spaces, such as SE(3), we consider both the position and orientation of the robot and redefine the state variable, \mathbf{x} , accordingly. We pick the *snake motion*—an overly complex and nonlinear motion presented in [7] and 2 different motions in the handwriting dataset, namely *Sine*- and *G-shaped* motions, to consolidate our higher dimensional performances. We use synthetic demonstrations for orientation, as the demonstrations only comprise translation and linear velocity. After training SNDS on this data, the policy generates trajectories in SE(3) that guide the end-effector toward the desired pose.

In Figure 7, we compare the performance of our method against the baselines trained with the Messy Snake data in SE(3). We utilize the DTW metric to evaluate the policy rollouts against the original trajectories. Next, we deploy the policy trained with SNDS on the simulated arm in PyBullet. Figure 8 offers an overview of policy deployment in simulation trained on the Messy Snake data.

Finally, we deploy the learned policy directly on a similar real-world manipulator platform. Recordings of experiments conducted with the Kinova Jaco2 arm are part of the supplementary video. Upon viewing the footage, it is evident that the robot reproduces the same motion as the simulation. The smooth sim-to-real transfer is made possible by the presence of Lyapunov stability conditions. The learned policy is robust to model error, since the Lyapunov condition ensures that the

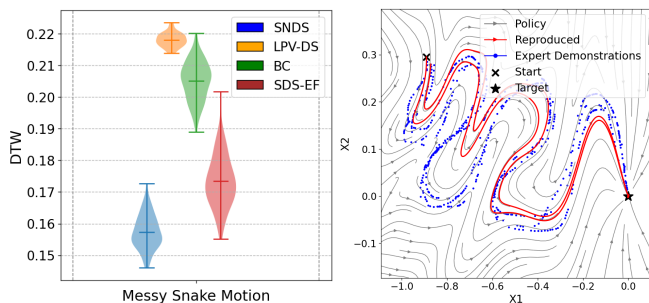


Fig. 7. DTW between policy rollouts and expert’s data for the complex Messy Snake demonstrations in SE(3) (left). SNDS can better extend to higher state space dimensions compared to former baselines. The trained SNDS policy over the expert demonstrations is depicted on the right.

robot consistently progresses toward the target.

V. DISCUSSION

The experiments demonstrate that employing SNDS to model a policy results in improved predictability and safety when recovering from states far beyond those covered in the original demonstrations. SNDS proves to be more effective by utilizing expressive and scalable neural architecture in comparison to the state of the art. Consequently, SNDS can acquire highly nonlinear stable policies in larger state spaces with reduced computational overhead. Our approach also reduces the restrictive assumptions regarding the class of Lyapunov functions to convexity, while the adopted projection technique ensures a smooth training process on a variety of expert data, as evident by the experiments.

While SNDS offers global convergence and predictability, there remains a concern regarding physically impossible trajectories for robots to replicate, particularly when considering manipulators and their joint or torque limits. Future enhancements could address this by incorporating safe regions using control barrier functions. Moreover, it is worth noting that modeling the Lyapunov candidate as a strictly convex function is not an essential requirement. Future research could explore invertible transformations to ease this restriction. Additionally, integrating obstacle avoidance and joint constraints into our formulation are promising paths to explore beyond the current scope. More ambitious extensions might delve into the implications of utilizing SNDS policies in reinforcement learning or training and deploying stable policies on legged robots, especially for gait control.

VI. CONCLUSION

We outlined a training process to learn expressive neural policies through imitating expert demonstrations. We effectively enforce global stability throughout the training process by adhering to the conditions of Lyapunov stability theory. This ensures that the resulting policy reliably converges to a predetermined target, regardless of initial conditions, velocity and time variations, or unexpected perturbations encountered during deployment. Our theoretical findings are accompanied by simulation and real-world benchmarking against best-performing methods in the field.




Fig. 8. Policy deployment in simulation (top) and real-world Kinova Jaco2 arm (bottom) scenarios. The policy is trained on a single Messy Snake demonstration. The simulation rollouts (red) are close to the expert’s demonstration (blue), and are similar to the rollouts in Figure 7.

APPENDIX

Network architecture and hyperparameters. Our approach employs a 4-hidden layer feed-forward neural network for policy, with Leaky ReLU activations and layer sizes of 2-256-256-128-128-2. For the Lyapunov function, we utilize an ICNN architecture [22] with 3 hidden layers of 2-128-128-128-1 nodes along with Leaky ReLU and Soft plus activations as required by the original design. We set $\gamma_i = \frac{1}{2^i}$ and $N_w = 2$ to limit the cost function’s horizon. Other parameters, e.g., learning rate and batch size, are selected and explained in our code repository. We use grid-search to pick the hyperparameters if computationally feasible.

Datasets specifications. The Handwriting Dataset [24] contains 30 sets of 2D handwriting motions recorded from a Tablet-PC via user input. There are seven demonstrations per motion, starting from slightly different initial positions but ending at $(0, 0)$. Each demonstration encompasses position (2×1000) as learning features and velocity (2×1000) as labels. We pick the following representative set of motion for our experiments: *Sine, P, Worm, G, N, DBLine, C, Angle*. Note that the *Messy Snake* motion is not a part of this dataset. The Messy Snake data, gathered in [7], comprises three demonstrations with various sample sizes per demonstration, totaling to (2×1592) , and we sample a single demonstration for policy learning and deployment.

Computational resources. We utilize a single machine with NVIDIA GeForce RTX 4050 GPU, Intel Core i7-13620H CPU, and 32 GB DDR4 RAM.

Codebase and reproduction. Consult `README.md` in our GitHub repository github.com/aminabyaneh/stable-imitation-policy to access SNDS experiments.

ACKNOWLEDGMENT

We express our heartfelt gratitude to Anya Forestell for her invaluable hardware support during our experiments.

REFERENCES

- [1] S. Schaal, "Is imitation learning the route to humanoid robots?" *Trends in cognitive sciences*, vol. 3, no. 6, pp. 233–242, 1999.
- [2] A. Hussein, M. M. Gaber, E. Elyan, and C. Jayne, "Imitation learning: A survey of learning methods," *ACM Computing Surveys (CSUR)*, vol. 50, no. 2, pp. 1–35, 2017.
- [3] N. Figueroa and A. Billard, "Locally active globally stable dynamical systems: Theory, learning, and experiments," *The International Journal of Robotics Research*, vol. 41, no. 3, pp. 312–347, 2022.
- [4] M. Hersch, F. Guenter, S. Calinon, and A. Billard, "Dynamical system modulation for robot learning via kinesthetic demonstrations," *IEEE Transactions on Robotics*, vol. 24, no. 6, pp. 1463–1467, 2008.
- [5] S. M. Khansari-Zadeh and A. Billard, "Learning stable nonlinear dynamical systems with gaussian mixture models," *IEEE Transactions on Robotics*, vol. 27, no. 5, pp. 943–957, 2011.
- [6] S. M. Khansari-Zadeh and A. Billard, "Learning control Lyapunov function to ensure stability of dynamical system-based robot reaching motions," *Robotics and Autonomous Systems*, vol. 62, no. 6, pp. 752–765, 2014.
- [7] N. Figueroa and A. Billard, "A physically-consistent bayesian non-parametric mixture model for dynamical system learning," in *2nd Annual Conference on Robot Learning*, 2018, pp. 927–946.
- [8] K. Neumann and J. J. Steil, "Learning robot motions with stable dynamical systems under diffeomorphic transformations," *Robotics and Autonomous Systems*, vol. 70, pp. 1–15, 2015.
- [9] H. Ravichandar, I. Salehi, and A. Dani, "Learning partially contracting dynamical systems from demonstrations," in *Conference on Robot Learning*. PMLR, 2017, pp. 369–378.
- [10] F. Khadivar, I. Lauzana, and A. Billard, "Learning dynamical systems with bifurcations," *Robotics and Autonomous Systems*, vol. 136, p. 103700, 2021.
- [11] A. Abyaneh and H.-C. Lin, "Learning Lyapunov-stable polynomial dynamical systems through imitation," in *7th Annual Conference on Robot Learning*, 2023.
- [12] Y.-C. Chang, N. Roohi, and S. Gao, "Neural Lyapunov control," *Advances in neural information processing systems*, vol. 32, 2019.
- [13] M. A. Rana, A. Li, D. Fox, B. Boots, F. Ramos, and N. Ratliff, "Euclideanizing flows: Diffeomorphic reduction for learning stable dynamical systems," in *Learning for Dynamics and Control*. PMLR, 2020, pp. 630–639.
- [14] J. Zhang, H. B. Mohammadi, and L. Rozo, "Learning riemannian stable dynamical systems via diffeomorphisms," in *6th Annual Conference on Robot Learning*, 2022.
- [15] J. Ho and S. Ermon, "Generative adversarial imitation learning," *Advances in neural information processing systems*, vol. 29, p. 4572–4580, 2016.
- [16] J. Fu, K. Luo, and S. Levine, "Learning robust rewards with adversarial inverse reinforcement learning," in *International Conference on Learning Representations*, 2018.
- [17] H. Dai, B. Landry, L. Yang, M. Pavone, and R. Tedrake, "Lyapunov-stable neural-network control," in *Robotics: Science and Systems*, 2021.
- [18] A. Coulombe and H.-C. Lin, "Generating stable and collision-free policies through Lyapunov function learning," in *International Conference on Robotics and Automation*, 2023, pp. 3037–3043.
- [19] P. Abbeel and A. Y. Ng, "Apprenticeship learning via inverse reinforcement learning," in *International conference on Machine learning*, 2004, p. 1.
- [20] B. D. Ziebart, A. L. Maas, J. A. Bagnell, A. K. Dey, *et al.*, "Maximum entropy inverse reinforcement learning," in *Association for the Advancement of Artificial Intelligence*, vol. 8, 2008, pp. 1433–1438.
- [21] R. L. Devaney, *An introduction to chaotic dynamical systems*. CRC press, 2021.
- [22] B. Amos, L. Xu, and J. Z. Kolter, "Input convex neural networks," in *International Conference on Machine Learning*. PMLR, 2017, pp. 146–155.
- [23] J. Z. Kolter and G. Manek, "Learning stable deep dynamics models," *Advances in neural information processing systems*, vol. 32, 2019.
- [24] S. M. Khansari-Zadeh and A. Billard, "LASA Handwriting Dataset," https://cs.stanford.edu/people/khansari/download.html#SEDS_reference, 2011.
- [25] M. Bruveris, "Optimal reparametrizations in the square root velocity framework," *SIAM Journal on Mathematical Analysis*, vol. 48, no. 6, pp. 4335–4354, 2016.
- [26] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *International Conference on Learning Representations*, 2015.
- [27] D. A. Pomerleau, "Alvinn: An autonomous land vehicle in a neural network," in *Advances in neural information processing systems*, vol. 1, 1988, pp. 305–313.
- [28] F. Torabi, G. Warnell, and P. Stone, "Behavioral cloning from observation," in *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, 2018, pp. 4950–4957.
- [29] E. Keogh and C. A. Ratanamahatana, "Exact indexing of dynamic time warping," *Knowledge and information systems*, vol. 7, pp. 358–386, 2005.