

LILAD: Learning In-context Lyapunov-stable Adaptive Dynamics Models

Amit Jena¹, Na Li², Le Xie²

¹ Department of Electrical and Computer Engineering, Texas A&M University, USA

² Harvard John A. Paulson School of Engineering and Applied Sciences, Harvard University, USA
 amit_jena@tamu.edu, nali@seas.harvard.edu, xie@seas.harvard.edu

Abstract

System identification in control theory aims to approximate dynamical systems from trajectory data. While neural networks have demonstrated strong predictive accuracy, they often fail to preserve critical physical properties such as stability and typically assume stationary dynamics, limiting their applicability under distribution shifts. Existing approaches generally address either stability or adaptability in isolation, lacking a unified framework that ensures both. We propose LILAD (Learning In-Context Lyapunov-stable Adaptive Dynamics), a novel framework for system identification that jointly guarantees adaptability and stability. LILAD simultaneously learns a dynamics model and a Lyapunov function through in-context learning (ICL), explicitly accounting for parametric uncertainty. Trained across a diverse set of tasks, LILAD produces a stability-aware, adaptive dynamics model alongside an adaptive Lyapunov certificate. At test time, both components adapt to a new system instance using a short trajectory prompt, which enables fast generalization. To rigorously ensure stability, LILAD also computes a state-dependent attenuator that enforces a sufficient decrease condition on the Lyapunov function for any state in the new system instance. This mechanism extends stability guarantees even under out-of-distribution and out-of-task scenarios. We evaluate LILAD on benchmark autonomous systems and demonstrate that it outperforms adaptive, robust, and non-adaptive baselines in predictive accuracy.

Introduction

Identifying dynamical systems from trajectory data is a foundational problem with wide-ranging applications in power systems, biological systems, and beyond. Despite the extensive literature on linear system identification, many real-world systems are inherently nonlinear and lie beyond the scope of linear models. Moreover, learning a linear approximation from data generated by a nonlinear system introduces modeling errors, often resulting in suboptimal performance. Many research works attempt to alleviate this challenge by employing nonlinear basis functions, such as sinusoidal series (Korenberg 1989) or sum-of-squares representations (Li, Dong, and Wang 2018), as well as iterative optimization methods like Monte Carlo techniques (Schön et al. 2015) and genetic algorithms (Chang 2007). However, these approaches rely on restrictive assumptions such as polynomial

structure or sufficient smoothness of dynamics. Such assumptions limit their applicability, and may lead to a failure in capturing the complex, high-frequency transient behavior exhibited by many real-world systems.

Neural networks offer a flexible system identification framework due to their expressive representational capacity, enabling modeling of a broader class of systems. However, due to their black box nature, the key physical properties of the underlying dynamics such as stability, monotonicity, dissipativity don't reflect in the learned models. Standard training on state-next-state pairs typically prioritizes accuracy over preserving these essential system characteristics, thereby inhibiting the deployability of such learned models in safety-critical applications.

A recent body of work has focused on embedding physical properties into learned models, aiming to create trustworthy digital twins of the underlying systems. One such approach (Kolter and Manek 2019) enforces stability by jointly training a dynamics model with a neural Lyapunov function. Another notable method (Kang et al. 2021) induces Lyapunov stability in neural ordinary differential equations (ODEs), making the resulting models robust to adversarial perturbations. Despite their strong empirical performance and theoretical foundations, these approaches suffer from a key limitation: the assumption of stationarity in the underlying system. However, many practical systems don't comply with this assumption, exhibiting time-varying dynamics or stochasticity in system parameters. For example, the topology of a power distribution network is often altered during the system operation for different objectives such as fault isolation or load balancing, leading to a time-varying dynamics. Another example is the seasonal variability in solar irradiance, which introduces parametric uncertainty in PV-inverter-integrated power systems. Among all physical properties, *stability* holds the most critical importance for ensuring reliable system operation. However, existing works overlook the possibility of non-stationary dynamics. A common remedy like retraining a stable dynamics model from scratch whenever the system changes, requires a new training dataset and significant additional training time, rendering it impractical for time-sensitive operations. Another parallel line of research employs adaptive learning methods such as meta-learning (Park et al. 2022) or in-context learning (Forgione, Pura, and Piga 2023) to model non-stationary dynamics under mild assumptions. However,

similar to standard neural network approaches, these methods aim to minimize squared loss errors without ensuring the stability of the learned dynamics. As a result, *learning stability-preserving dynamics models that remain effective under mild constraints in non-stationary settings remains an open and challenging problem.*

In this paper, we study the problem of learning stability-preserving adaptive dynamics models. We propose **LILAD** (Learning In-context Lyapunov-stable Adaptive Dynamics Models), a novel framework that leverages the in-context learning (ICL) capability of a large language model (LLM) to approximate nonstationary dynamics, while enforcing stability through Lyapunov-based conditions. We assume parametric stochasticity in the system, with access to a data pool comprising multi-task trajectory datasets. Each dataset is generated by instantiating the system with a specific parameter sample. During the training of LILAD, we employ the ICL procedure to jointly train a dynamics model and a Lyapunov function that satisfies global exponential stability across the multi-task data pool. Specifically, we adopt an adversarial training framework that alternates between updating the dynamics model and the Lyapunov function to promote joint and efficient learning. At evaluation, LILAD receives a small test-time trajectory dataset as a prompt corresponding to a new, unseen parameter sample, and produces a stable dynamics in a zero-shot manner. Furthermore, using a bisection-based root-finding method, we compute an attenuator γ using the trained Lyapunov and dynamics models, which strictly enforces stability in the learned dynamics model. We validate LILAD against baseline methods on several benchmark autonomous systems, including a high-dimensional example to demonstrate the adaptability and stability performance of our approach. Our codes are available at <https://github.com/amitjena1992/LILAD>.

Related Work

System Identification: System Identification encompasses a broad range of algorithms, all aimed at approximating models that accurately predict the future state of a system given its current state. Traditional modeling approaches include linear models (Mehra 1974), as well as representation using basis functions like sinusoidal series (Korenberg 1989), sum-of-squares polynomials (Li, Dong, and Wang 2018), kernels like quadratic and Gaussian functions (Greblicki and Pawlak 2003). Another class of works focuses on block-oriented methods, including Wiener models (Hagenblad 1999), Hammerstein models (Chaoui et al. 2005), and hybrid Wiener–Hammerstein structures (Sun et al. 2022), which combine linear blocks and static nonlinearity to approximate nonlinear systems. In parallel, a different line of study utilizes iterative algorithms such as genetic algorithms (Chang 2007) and Monte Carlo methods (Schön et al. 2015), as well as Bayesian approaches like Gaussian process regression (Bijl et al. 2017) to approximate nonlinear dynamics in a data-efficient and uncertainty-aware manner. More recent advances advocate for neural network-based techniques that are highly expressive and free from restrictive model assumptions. Examples include convolutional neural networks (CNNs) (Wu and Jahanshahi 2019), Neural ODEs

(Quaglino et al. 2019), and physics-informed neural networks (PINNs) (Stiasny, Misyris, and Chatzivasileiadis 2021).

Stability-preserving System Identification: Identifying stable systems is of utmost priority, particularly when the learned model is intended for deployment in a safety-critical system. These concerns have laid the foundation for a dedicated section of work. In this setting, various methods have been proposed that rigorously guarantee stability in the approximated dynamics. Classical approaches include stable spline kernels (Pillonetto and De Nicolao 2010), Lyapunov-based frameworks (Lyashevskiy and Abel 1994), subspace methods (Maciejowski 1995), convex parameterizations for stable systems (Tobenkin et al. 2010), and approximation of stable positive systems (Umenberger and Manchester 2016). With the advent of deep learning, this line of research has advanced with learning-based techniques like Lyapunov-based neural network methods (Kolter and Manek 2019; Lawrence et al. 2020), stable neural ODEs (Kang et al. 2021), stability-aware deep generative models (Urain et al. 2020), and Lyapunov-stable deep equilibrium models (Chu et al. 2024).

Adaptive and Robust System Identification: The presence of stochasticity in a dynamical system necessitates adaptive and robust methods for reliable operation and control. This challenge has been addressed through both classical and modern approaches. Classical adaptive linear methods include least mean square (LMS)-based techniques (Chen, Gu, and Hero 2009), adaptive Wiener models (Ogunfunmi 2007), and robust error-domain model falsification (EDMF) methods (Pasquier and Smith 2015). More recently, learning-based approaches have gained traction, including meta-learning (Park et al. 2022) and in-context learning (ICL) (Forgione, Pura, and Piga 2023; Jena et al. 2025), which offer greater flexibility in handling complex, nonstationary dynamics.

Despite progress in learning-based stable and adaptive modeling separately, no prior work has unified both objectives. To the best of our knowledge, this is the first neural network-based approach to address stable adaptive system identification, demonstrating improved performance over standard adaptive baselines.

Preliminaries

Stability of Discrete-Time Dynamical Systems

We summarize the fundamental stability results for a generic discrete-time deterministic system of the form:

$$x_{k+1} = f(x_k), \quad (1)$$

where $x_k \in X \subset \mathbb{R}^d$ denotes the system state at discrete time k , and $f(\cdot) : \mathbb{R}^d \rightarrow \mathbb{R}^d$ represents the dynamics. The equilibrium point of the system (1) is assumed to be $x = 0$.

Definition 1 (Global exponential Stability (GES) (Grujić and Šiljak 1974)). The equilibrium $x = 0$ is said to be globally exponentially stable if and only if there exist constants $m > 1$, $\alpha > 0$ such that

$$\|x_k\|_2 \leq m \|x_0\|_2 e^{-\alpha k}, \forall k \in \mathbb{N}, \quad (2)$$

for any initial state $x_0 \in X$.

GES implies global asymptotic stability where $x_k \rightarrow 0$ as $k \rightarrow \infty$. Additionally, it indicates that the rate of decay is exponential in time with decay rate α .

Definition 2 (Lyapunov functions for discrete-time systems (Grujić and Šiljak 1974)). The equilibrium $x = 0$ is globally exponentially stable for system (1) if and only if there exists a scalar-valued function $V : \mathbb{R}^d \rightarrow \mathbb{R}$ with the following properties:

$$V(x) \text{ is continuous in } \mathbb{R}^d, \quad (3a)$$

$$c_1 \|x\|_2^2 \leq V(x) \leq c_2 \|x\|_2^2, \quad \forall x \in X, \quad (3b)$$

$$V(x_{k+1}) \leq \beta V(x_k), \quad \forall x_k \in X, \quad \forall k \in \{0, 1, \dots\} \quad (3c)$$

where $c_1 > 0$, $c_2 > 0$, $0 < \beta < 1$.

Condition (3b) implies that $V(x) > 0$ for all $x \in X \setminus \{0\}$ and $V(x) = 0$ if and only if $x = 0$, consistent with other variants of Lyapunov-based definitions, such as those in (Bof, Carli, and Schenato 2018). Additionally, many existing works (Lawrence et al. 2020) conventionally consider $\beta = 1 - \alpha$ where $\alpha \in (0, 1)$ is the decay rate of the dynamics as defined in Definition 1. The existence of a Lyapunov function satisfying the conditions in Definition 2 ensures that the system (1) is globally exponentially stable, as stated in Definition 1.

In-context Learning

We briefly summarize the key principles of ICL, the underlying basis of our proposed LILAD framework. Let \mathcal{T} be a family of tasks, and let \mathcal{P}_τ denote a meta-distribution across \mathcal{T} . Any task $\tau \in \mathcal{T}$ induces a task distribution $P_{Z|\tau}$ on the domain $Z = U \times V$, where U and V are the input space and output space, respectively. Given a random prompt

$$\mathcal{P}^\tau = \{u_1^\tau, v_1^\tau, u_2^\tau, v_2^\tau, \dots, u_n^\tau, v_n^\tau, u_{n+1}^\tau\} \quad (4)$$

which comprises n in-context input-output pairs and a query input u_{n+1}^τ , where each $(u_i^\tau, v_i^\tau) \sim P_{Z|\tau}$ i.i.d. for $i \in \{1, \dots, n\}$. Additionally, the query input and true output $(u_{n+1}^\tau, v_{n+1}^\tau) \sim P_{Z|\tau}$ are also drawn independently from the same distribution. In this setting, an ICL model $M(\cdot)$ aims to predict v_{n+1}^τ with an average error bounded by ε . Formally,

$$\mathbb{E}_{\mathcal{P}_\tau} \mathbb{E}_{P_{Z|\tau}^{n+1}} [\ell(M(\mathcal{P}^\tau), v_{n+1}^\tau)] \leq \varepsilon, \quad (5)$$

where $\ell(\cdot, \cdot)$ is an appropriate loss function. In this process, the model $M(\cdot)$ adapts to any task τ solely through the prompt \mathcal{P}_τ , without requiring any gradient updates.

ICL of Adaptive and Stable Dynamics

We consider a discrete-time autonomous system of the form

$$x_{k+1} = f_\vartheta(x_k), \quad (6)$$

where $x_k \in X \subset \mathbb{R}^d$ denotes the fully observed time-evolving state of the system, $f_\vartheta : \mathbb{R}^d \rightarrow \mathbb{R}^d$ represents the parametrized dynamics, and $\vartheta \in \mathbb{R}^\vartheta$ is a stochastic system parameter governed by an unknown probability distribution \mathcal{P}_ϑ . Although the distribution \mathcal{P}_ϑ is explicitly unknown, its variability is experienced through a finite set of parameter samples $\mathcal{D}^\vartheta := \{\tilde{\vartheta}_1, \tilde{\vartheta}_2, \dots, \tilde{\vartheta}_M\}$, where each $\tilde{\vartheta}_i$ leads to a

distinct deterministic dynamics $f_{\tilde{\vartheta}_i}(\cdot)$. We assume that the origin is the equilibrium point of each such dynamical system, i.e., $f_{\tilde{\vartheta}_i}(0) = 0$ holds for every $\tilde{\vartheta}_i \in \mathcal{D}^\vartheta$. Next, we make the following assumption regarding the stability of the dynamics, which is fundamental to our formulation.

Assumption 1 (GES). Let $\tilde{\vartheta}_i \in \mathcal{D}^\vartheta$ be a parameter sample, and let $f_{\tilde{\vartheta}_i}(\cdot)$ be its corresponding induced dynamics. Then for some $m > 1$ and $\alpha > 0$, the dynamics $x_{k+1} = f_{\tilde{\vartheta}_i}(x_k)$ satisfies GES condition as defined in Definition 1.

Assumption 1 ensures that any induced dynamics $f_{\tilde{\vartheta}_i}(\cdot)$ is globally exponentially stable, meaning its trajectories converge to the equilibrium at an exponential rate for any initial condition. This type of assumption is widely used in the stability analysis literature (Grujić and Šiljak 1974; Geiselhart et al. 2014) in various forms. Our problem formulation is developed under this assumption.

Approximating the nonstationary dynamics as stated in (6) from trajectory data under parametric uncertainty is challenging, especially when stability must be preserved. We address this by proposing an adaptive ICL-based method that ensures stable dynamics. In our set-up, we assume access to a training data pool $\mathcal{D}^{\text{train}} := \{\mathcal{D}^i\}_{i=1}^M = \{(x_{i,j}, f_{\tilde{\vartheta}_i}(x_{i,j}))\}_{j=1}^N\}_{i=1}^M$, consisting of M datasets. Each dataset \mathcal{D}^i corresponds to a distinct dynamics $f_{\tilde{\vartheta}_i}(\cdot)$ induced by $\tilde{\vartheta}_i \in \mathcal{D}^\vartheta$, and contains N state-next-state pairs, generated from rollouts initialized at random initial states. We note that the state-next-state pairs have been randomly shuffled to remove the temporal dependencies between samples, resulting in i.i.d. datasets suitable for standard learning algorithms. Thus, $x_{i,j}$ and $x_{i,j+1}$ are not sequential states but are independently sampled from the dataset.

We present LILAD, an ICL-based framework that jointly learns adaptive dynamics and Lyapunov models. LILAD training proceeds by alternating between two phases: (i) updating the Lyapunov model using a *Lyapunov loss* computed with the current fixed dynamics model, and (ii) updating the dynamics model using a combination of standard MSE and a *dynamics correction loss* derived from the fixed Lyapunov model. This process continues until both models converge, resulting in a stability-aware dynamics model and a corresponding Lyapunov function.

Both models jointly satisfy the Lyapunov stability constraints (as defined in Definition 2) for most samples in the training data pool. However, like many neural network-based methods, the *out-of-task* and *out-of-sample* performances of LILAD remain uncertain. We note that due to the design of our Lyapunov model architecture, the first two conditions (3a) and (3b) are inherently satisfied. Consequently, we focus on ensuring the satisfaction of the final condition, (3c), which enforces exponential stability. To this end, we utilize a bisection-based root-finding method to compute an attenuation factor $\gamma \in [0, 1)$ that guarantees Lyapunov decrease across out-of-task and out-of-sample scenarios. The key idea is to sufficiently attenuate the dynamics model to ensure a suitable decrease in the Lyapunov value between two consecutive states, thereby satisfying the exponential stability condition. It is important to note that γ is not computed globally; instead, it is determined separately for each state,

making it inherently state-dependent. In what follows, first we describe the LILAD set-up in detail, and then proceed to the training and evaluation procedures of LILAD.

LILAD Framework Set-up

We begin by introducing the following notations for our adaptive system identification framework:

$$\mathcal{P}^i := \{x_{i,1}, f_{\hat{\vartheta}_i}(x_{i,1}), \dots, x_{i,n}, f_{\hat{\vartheta}_i}(x_{i,n}), x_{i,n+1}\} \quad (7a)$$

$$\mathcal{P}_j^i := \{x_{i,1}, f_{\hat{\vartheta}_i}(x_{i,1}), \dots, x_{i,j}, f_{\hat{\vartheta}_i}(x_{i,j}), x_{i,j+1}\}, \quad (7b)$$

$$\mathcal{C}^i = \mathcal{P}^i \setminus \{x_{i,n+1}\}, \quad \mathcal{C}_j^i = \mathcal{P}_j^i \setminus \{x_{i,j+1}\}, \quad (7c)$$

where \mathcal{P}^i denotes the full prompt for the i -th task, consisting of n in-context state–next-state pairs \mathcal{C}^i and a query $x_{i,n+1}$. This prompt is constructed by randomly selecting n pairs from the N available state–next-state pairs associated with i -th task. Similarly, \mathcal{P}_j^i is the (i, j) -th prompt prefix, comprising the first j in-context pairs \mathcal{C}_j^i and query $x_{i,j+1}$. We refer to \mathcal{C}^i and \mathcal{C}_j^i as i -th context and (i, j) -th context prefix, respectively. In this setting, the proposed ICL-based dynamics model G_θ takes any prompt prefix \mathcal{P}_j^i as input and predicts $f_{\hat{\vartheta}_i}(x_{i,j+1})$. In parallel, the ICL-based Lyapunov model V_ϕ also receives \mathcal{P}_j^i and outputs the Lyapunov value corresponding to $x_{i,j+1}$. We present these predictions as follows:

$$\hat{f}(x_{i,j+1}) := G_\theta(\mathcal{P}_j^i) = G_\theta(x_{i,j+1} | \mathcal{C}_j^i), \quad (8a)$$

$$\hat{V}(x_{i,j+1}) := V_\phi(\mathcal{P}_j^i) = V_\phi(x_{i,j+1} | \mathcal{C}_j^i), \quad (8b)$$

where $\hat{f}(x_{i,j+1})$ and $\hat{V}(x_{i,j+1})$ denote the predicted dynamics and Lyapunov values at $x_{i,j+1}$ conditioned on the context \mathcal{C}_j^i . We note that the context prefix \mathcal{C}_0^i is empty for all cases. However, this does not inhibit prediction, as the ICL models can still generate outputs based on the query input alone, without any contextual information.

The architectures for G_θ and V_ϕ follow a common structure comprising a linear input layer, a GPT-2 transformer block, and a linear output layer. In both models, the input layer tokenizes the prompt and passes it to the transformer block, which generates a context-aware embedding for a prediction on the query. Then, this prediction is passed through the output layer to produce the final model output. While standard transformer architectures ensure that the first Lyapunov constraint (3a) holds for any general query conditioned on a fixed context, the second condition (3b) isn't inherently satisfied. We make the following architectural modifications to overcome this challenge.

Enforcing Semi-definiteness via Output Warping We employ a scaled tanh activation after the output layer to constrain the Lyapunov output within the range $[-c, c]$ for some $c > 0$, which improves the numerical tractability of integrating V_ϕ into the training pipeline. To further enforce the semi-definiteness of V_ϕ , we adopt a smoothed ReLU activation inspired by (Kolter and Manek 2019), defined as:

$$\sigma(x) = \begin{cases} 0 & \text{if } x \leq 0 \\ x^2/2\delta & \text{if } 0 < x < \delta \\ x - \delta/2 & \text{otherwise,} \end{cases} \quad (9)$$

where $\delta > 0$ is a tunable parameter regulating the smoothness of the function. Then, the following output warping scheme is applied to ensure that V_ϕ remains positive semi-definite:

$$V_\phi(x_{i,j+1} | \mathcal{C}_j^i) = \sigma(c \cdot \tanh(V_\phi^{\text{raw}}(x_{i,j+1} | \mathcal{C}_j^i)) - c \cdot \tanh(V_\phi^{\text{raw}}(0 | \mathcal{C}_j^i))) + \epsilon \|x_{i,j+1}\|^2, \quad (10)$$

where $V_\phi^{\text{raw}}(\cdot | \mathcal{C}_j^i)$ denotes unwrapped Lyapunov output, and $\epsilon > 0$ ensures a unique minima at $x_{i,j+1} = 0$. Now, it's straightforward to prove that with the positive semi-definiteness and boundedness constraints imposed, $V_\phi(x_{i,j+1} | \mathcal{C}_j^i)$ satisfies the second Lyapunov condition (3b) for any $x_{i,j+1} \in X$. In the remainder of this section, we separately describe the training and evaluation procedures for LILAD.

LILAD Training

We adopt an adversarial training procedure consisting of two stages. In the first stage, the dynamics model G_θ is frozen while V_ϕ is trained. In the second phase, V_ϕ is fixed and G_θ is updated. This alternating process progresses until both models converge in terms of their respective loss functions. The dynamics and Lyapunov losses, represented by $\mathcal{L}^{\text{Dyn}}(\theta)$ and $\mathcal{L}^{\text{Lyap}}(\phi)$, are defined as follows:

$$\begin{aligned} \mathcal{L}^{\text{Dyn}}(\theta) = & \frac{1}{M(n+1)} \sum_{i=1}^M \sum_{j=0}^n \|G_\theta(x_{i,j+1} | \mathcal{C}_j^i) \\ & - f_{\hat{\vartheta}_i}(x_{i,j+1})\| + \lambda \max \{V_\phi(G_\theta(x_{i,j+1} | \mathcal{C}_j^i) | \mathcal{C}_j^i) \\ & - \beta V_\phi(x_{i,j+1} | \mathcal{C}_j^i), 0\}, \end{aligned} \quad (11a)$$

$$\begin{aligned} \mathcal{L}^{\text{Lyap}}(\phi) = & \frac{1}{M(n+1)} \sum_{i=1}^M \sum_{j=0}^n \max \{ \\ & V_\phi(G_\theta(x_{i,j+1} | \mathcal{C}_j^i) | \mathcal{C}_j^i) - \beta V_\phi(x_{i,j+1} | \mathcal{C}_j^i), 0\}, \end{aligned} \quad (11b)$$

where $\lambda > 0$ is a small regularization constant, and β is the exponential stability coefficient, as defined in (3c). For the dynamics loss $\mathcal{L}^{\text{Dyn}}(\theta)$, the first term inside the summation is a standard MSE between the context-conditioned predicted next state and the true next state. The second term penalizes violations of the Lyapunov condition (3c), serving as a dynamics correction term while V_ϕ is held fixed. Similarly, for $\mathcal{L}^{\text{Lyap}}(\phi)$, the dynamics model G_θ is frozen and the loss captures the extent to which the current V_ϕ violates the Lyapunov condition (3c). Thus, both models contribute to the evaluation of both loss functions, enabling a joint and efficient training process. Upon convergence, the training algorithm yields a stability-aware dynamics model G_{θ^*} and a Lyapunov function V_{ϕ^*} , both optimized with respect to their respective loss functions. The complete LILAD training procedure is summarized in Algorithm 1.

LILAD Evaluation and Stability Enforcement

The optimal models G_{θ^*} and V_{ϕ^*} obtained from LILAD training jointly satisfy Lyapunov constraints for the majority of samples in the training prompts $\{\mathcal{P}_i\}_{i=1}^M$. However, deploying such models in safety-critical systems demands rigorous stability guarantees, which go beyond empirical

performance and require stronger, verifiable assurances. We address this bottleneck by presenting a bisection-based root-finding method to compute a state-dependent attenuator γ , which rigorously enforces Lyapunov stability for arbitrary states, thereby generalizing to both out-of-task and out-of-sample scenarios. Our approach builds on the technique introduced in (Lawrence et al. 2020), which was originally developed for stationary dynamics. We extend this idea to the non-stationary setting considered in this work. We note that, contrary to many prior works (Kolter and Manek 2019; Chu et al. 2024), our method does not assume input-output convexity of the Lyapunov model V_ϕ , a property that is difficult to verify for transformer-based LLMs.

The strategy followed in this step is to first check if the context-conditioned next-state prediction by the optimal dynamics model satisfies the third Lyapunov condition (3c). Specifically, at any arbitrary query state $x \in X$, given a new test-time context prefix \mathcal{C}_j^{M+1} corresponding to a new $(M+1)$ -th task, we verify if $V_{\phi^*}(G_{\theta^*}(x|\mathcal{C}_j^{M+1})|\mathcal{C}_j^{M+1}) - \beta V_{\phi^*}(x|\mathcal{C}_j^{M+1}) < 0$. If this condition is satisfied at x , we set $\gamma = 1$. Otherwise, we compute a suitable γ as follows.

Find the largest $\gamma \in [0, 1)$ such that

$$V_\phi^*(\gamma \cdot G_\theta^*(x|\mathcal{C}_j^{M+1})|\mathcal{C}_j^{M+1}) - \beta V_\phi^*(x|\mathcal{C}_j^{M+1}) = 0. \quad (12)$$

This results in a one-dimensional root-finding problem, solvable with bisection, that ensures exponential stability of the modulated dynamics $\gamma \cdot G_\theta^*$. We utilize the intermediate value theorem (IVT) to guarantee the existence of a valid γ satisfying (12). The detailed proof of this claim has been deferred to the Appendix. It is important to note that γ is state-dependent, i.e., it varies with each query state and is not shared across tasks or samples. Hence, we denote it as $\gamma(x)$. Finally, the stability-constrained optimal dynamics model is represented as $\gamma(x) \cdot G_\theta^*(x|\mathcal{C}_j^{M+1})$, where

$$\gamma(x) = \begin{cases} 1 & \text{if } V_{\phi^*}(G_{\theta^*}(x|\mathcal{C}_j^{M+1})|\mathcal{C}_j^{M+1}) \\ & -\beta V_{\phi^*}(x|\mathcal{C}_j^{M+1}) < 0 \\ \text{solution of (12)} & \text{otherwise.} \end{cases} \quad (13)$$

Experiments

The effectiveness of our proposed LILAD is demonstrated through comparisons with standard baselines across multiple benchmark systems. Specifically, we consider five autonomous systems, each of which is inherently stable and either does not require external control or is already equipped with a stabilizing controller, such as pre-defined droop controllers in the case of the networked microgrid. As explained in earlier sections, we focus on the system identification task under parametric uncertainty. The benchmark systems used in our simulations include: (i) a one-link pendulum, (ii) a two-link pendulum, (iii) a networked microgrid, (iv) an SEIR epidemic model, and (v) a high-dimensional surrogate model of a reaction-diffusion partial differential equation (PDE).

Baselines

We consider four baseline methods for performance comparison. Our choice of baselines includes two adaptive al-

Algorithm 1: LILAD Training Algorithm

- 1: **Input:** Offline prompts $\{\mathcal{P}_i\}_{i=1}^M$, with each \mathcal{P}_i comprising n state-next-state pairs; switching interval K^{sw} .
 - 2: **Initialize:** Dynamics and Lyapunov model parameters $\theta^{(0)}, \phi^{(0)}$
 - 3: **for** $k = 0, \dots, K - 1$ **do**
 - 4: Randomly sample a batch of prompt prefixes $\{\mathcal{P}_i^p\}_{i=1}^M$, where $p < n$.
 - 5: **if** $\lfloor (k/K^{\text{sw}}) \rfloor \bmod 2 = 0$ **then**
 - 6: Compute and update dynamics loss $\mathcal{L}^{\text{Dyn}}(\theta^{(k)})$ as described in (11a) over $\{\mathcal{P}_i^p\}_{i=1}^M$ while fixing $V_{\phi^{(k)}}$.
 - 7: **else**
 - 8: Compute and update Lyapunov loss $\mathcal{L}^{\text{Lyap}}(\phi^{(k)})$ as described in (11b) over $\{\mathcal{P}_i^p\}_{i=1}^M$ while freezing $G_{\theta^{(k)}}$.
 - 9: **end if**
 - 10: **end for**
 - 11: **Output:** $\theta^* = \theta^{(K)}, \phi^* = \phi^{(K)}$
-

gorithms, one robust algorithm, and a non-adaptive stable system identification method. The adaptive baselines consist of standard ICL and model-agnostic meta-learning (MAML) (Finn, Abbeel, and Levine 2017), both of which have been widely used in adaptive regression and non-linear system identification tasks (Forgione, Pura, and Piga 2023; Park et al. 2022). In case of ICL, after a full course of in-context training, the ICL model makes predictions on a general query state given a test-time context prefix \mathcal{C}_j^{M+1} , similar to our proposed LILAD framework. In contrast, MAML does not depend on the ordering of \mathcal{C}_j^{M+1} ; it treats the test-time context as an unordered set of examples and performs gradient-based fine-tuning to adapt to new tasks. For the robust baseline, we adopt a Conditional Value-at-Risk (CVaR)-based regression method (Rockafellar, Uryasev et al. 2000), which employs a risk-sensitive max-min formulation to learn a system identification model that performs well under worst-case scenarios. Specifically, we follow a neural network-based approach that, instead of minimizing the average loss, identifies the top k worst-case losses, and optimizes over their average value. This method doesn't require a test-time prompt as it directly optimizes in a risk-aware manner during training, yielding a fixed model that generalizes to uncertain or adverse conditions. Finally, we include a linear stable system identification approach (referred to as stable-linear) that enforces stability by constraining the eigenvalues of the learned state-transition matrix to lie in the unit circle. We note that this method serves as a non-adaptive and non-robust baseline, as it is directly computed on the test-time context prefix. Finally, under parametric uncertainty, each method (including LILAD) is evaluated on its ability to adapt to new test-time tasks while maintaining stability in the identified system.

Simulations

We use mean absolute error (MAE) and root mean squared error (RMSE) as performance metrics to assess the efficacy




Figure 1: *Performance Comparison of LILAD*. For the Simple Pendulum system with stochastic parameters (g, l, b), the proposed LILAD method adapts more effectively to test-time instances than ICL. (a) ICL yields a sub-optimal approximation of the test-time vector field, causing predicted trajectories to deviate beyond the specified boundary. (b) In contrast, LILAD achieves a more accurate approximation, with guaranteed convergence of the predicted trajectories to the equilibrium. (c) Ground truth: the true vector field and corresponding reference trajectories of the test-time system.

ENV.	DIM	METRIC	ICL	MAML	CVAR	STABLE-LINEAR	LILAD
SP	2	MAE	0.018 ± 0.068	0.023 ± 0.014	0.085 ± 0.022	0.065 ± 0.016	0.004 ± 0.020
		RMSE	0.022 ± 0.080	0.034 ± 0.024	0.17 ± 0.09	0.13 ± 0.021	0.011 ± 0.012
DP	4	MAE	0.039 ± 0.029	0.022 ± 0.010	0.12 ± 0.05	0.17 ± 0.16	0.011 ± 0.008
		RMSE	0.061 ± 0.028	0.032 ± 0.008	0.22 ± 0.12	0.42 ± 0.25	0.021 ± 0.018
MG	5	MAE	0.005 ± 0.010	0.007 ± 0.009	0.014 ± 0.028	0.011 ± 0.004	0.005 ± 0.002
		RMSE	0.009 ± 0.012	0.009 ± 0.013	0.042 ± 0.068	0.017 ± 0.004	0.012 ± 0.009
SEIR	8	MAE	0.022 ± 0.017	0.032 ± 0.027	0.077 ± 0.05	1.049 ± 0.818	0.017 ± 0.019
		RMSE	0.031 ± 0.011	0.043 ± 0.067	0.18 ± 0.13	1.617 ± 1.404	0.022 ± 0.029
PDE-SM	100	MAE	6.354 ± 0.050	–	–	–	0.060 ± 0.002
		RMSE	7.478 ± 0.039	–	–	–	0.249 ± 0.008

Table 1: Comparison of MAE and RMSE of all Methods Across Dynamical Systems

of each method. Specifically, we evaluate a system identification method by first rolling out trajectories using the model adapted to the test-time task and then computing the MAE and RMSE between these predicted trajectories and the ground-truth trajectories generated from the actual test-time dynamics. To account for parametric uncertainty during evaluation, we instantiate five test-time systems by drawing five random parameter samples. For each of these test-time systems, we fix four initial states, which are used to generate both the ground truth trajectories and predicted trajectories. Finally, for each system, we compute a total of twenty MAE and RMSE values, one for each trajectory. The mean and standard deviation of these MAE and RMSE values for each system are reported in Table 1. In this comparison table, we use shorthands SP, DP, MG, SEIR, PDE-SM to denote the simple pendulum, double pendulum, networked microgrids, SEIR, and the PDE surrogate system, respectively.

We consider a simple pendulum setting as our first system, and assume the pendulum’s length (l), damping coefficient (b), and acceleration due to gravity (g) to be stochastic parameters. The system states are angular displacement from the vertically downward position θ and angular velocity $\dot{\theta}$. We compute the MAE and RMSE metrics for each method

and observe that LILAD outperforms all baselines, including the stable-linear method that utilizes the test-time data directly, and is already at an advantage. Quantitatively, LILAD obtains $4.5\times$ lower mean MAE and $2\times$ lower mean RMSE compared to ICL, the next best-performing approach. We further compare LILAD and ICL visually in Fig. 1, which shows the phase plots of a representative test-time system instance for LILAD, ICL, and the ground-truth dynamics. Each plot includes two rollout trajectories from the same initial points. We notice that LILAD more accurately captures the underlying vector field of the actual dynamics compared to ICL. Furthermore, both trajectories generated by LILAD stay within the bounding box as the ground truth trajectories, whereas ICL results in a trajectory that deviates beyond these bounds. This indicates that LILAD provides superior approximation quality, not only offering better accuracy but also better approximating the system’s vector field. This gives our method a strong advantage for deployment in safety-critical settings. Among the rest of the baselines, all significantly perform worse than both ICL and LILAD, with the exception of MAML, which exceeds ICL but underperforms compared to LILAD.

Our second benchmark is a double pendulum setting,




Figure 2: LILAD scales effectively to a high-dimensional system derived from a discretized reaction-diffusion PDE. Its learned surrogate predicts trajectories that converge to the origin, matching the test-time ground truth. In contrast, ICL does not guarantee convergence to the origin. Each image shows the spatial temperature profile under a fixed test-time diffusion coefficient.

where the two masses m_1 and m_2 , the length of two links l_1 and l_2 , damping coefficients b_1 and b_2 are assumed to take stochastic values. The states are $\theta_1, \theta_2, \dot{\theta}_1$, and $\dot{\theta}_2$, that represent the angular displacements and velocities. We evaluate all methods using the MAE and RMSE and observe that LILAD consistently supersedes all baselines in both metrics, mimicking its strong performance in the simple pendulum system. Specifically, LILAD fetches $2\times$ lower mean MAE and $1.5\times$ lower mean RMSE than MAML, which is the next best method for this setting. The ICL falls behind both MAML and LILAD, but yields substantially better results than CVaR and stable-linear approaches.

The third benchmark is a network of microgrids equipped with local phase-angle control loops. The droop control governs the overall dynamics in conjunction with AC power flow constraints. The state comprises the modified voltage phase angles $[\Delta\delta_1, \dots, \Delta\delta_5]$, where each $\Delta\delta_j = \delta_j - \delta_j^*$ denotes the difference between the j -th phase angle and its setpoint (i.e., equilibrium value). We assume that the droop constants $[D_{\delta_1}, \dots, D_{\delta_5}]$ to be stochastic parameters. We evaluate all methods in terms of MAE and RMSE and find that LILAD matches ICL in mean MAE and slightly lags behind ICL in mean RMSE. This behavior is attributed to LILAD’s more accurate approximation of the steady-state part, despite a minor under-approximation of the transient part. Both LILAD and ICL outperform the remaining baselines in this setting.

The fourth system models the progress of a disease outbreak in a population. It divides the general population into four groups: susceptible (S), exposed (E), infectious (I), and removed (R). We consider a scenario involving two interacting populations undergoing active vaccination campaigns. The state vector includes $S_1, E_1, I_1, R_1, S_2, E_2, I_2, R_2$ that indicate the population density in each compartment across both populations. The self transmission rates (β_{11}, β_{22}), and recovery rates γ_1 and γ_2 are treated as stochastic parameters. Based on the computed MAE and RMSE values for all methods, we observe that LILAD outperforms all baselines, demonstrating its superior adaptability to diverse and het-

erogeneous environments. In sharp contrast, both CVaR and stable-linear provide the weakest performance.

Scalability of LILAD to High-dimensional Systems To demonstrate the scalability of LILAD, we take a stable reaction-diffusion PDE that governs the spatiotemporal evolution of temperature across a surface. The surface is discretized into a 10×10 grid, yielding 100 nodal temperatures and resulting in a dynamical system of dimension 100. We treat the diffusion coefficient α_{diff} as stochastic and evaluate LILAD and ICL across five test-time instances. Consistent with previous results, we find that ICL doesn’t always guarantee the convergence of the adapted model to the origin during test time. In contrast, the predicted trajectories by LILAD always reach equilibrium, thereby closely matching the test-time ground truth. We demonstrate this behavior in Figure 2, where a LILAD-based representative trajectory is visually compared and contrasted with ICL’s diverging prediction using a sequence of heatmaps. Furthermore, as reported in Table 1, LILAD achieves approximately 106 times lower mean MAE and 30 times lower RMSE than ICL, suggesting the much higher reliability of LILAD in high-dimensional non-stationary safety-critical applications. In this setting, we exclude MAML, CVaR, and stable-linear baselines due to their lack of scalability and consistently subpar performance on large-scale systems.

Conclusion

In this work, we proposed LILAD, a novel system identification framework that unifies stability and adaptability, two properties typically studied in isolation. Utilizing in-context learning (ICL), LILAD employs two transformer models to jointly learn a stability-aware adaptive dynamics model and a valid adaptive Lyapunov function from a multi-task trajectory pool. During evaluation, both models adapt seamlessly to new dynamics using only contextual information in the form of test-time prompts. To further ensure stability, we incorporated a state-dependent attenuator γ , which guarantees sufficient Lyapunov decrease across predicted state transitions. We validated LILAD’s effectiveness across benchmark systems, demonstrating superior performance over adaptive, robust, and non-adaptive baselines. To the best of our knowledge, LILAD is the first neural network-based framework to achieve adaptive system identification while guaranteeing stability across diverse tasks.

In future work, we aim to extend LILAD to state estimation and particle filtering problems under parametric uncertainty, as well as explore adaptive control through ICL from multi-task data.

Acknowledgements

This work was supported in part by the NSF AI Institute (Award No. 2112085) and the Salata Institute for Climate and Sustainability at Harvard University.

The views expressed in this paper are the opinion of the authors and do not reflect the views of PJM Interconnection, L.L.C. or its Board of Managers of which Le Xie is a member.

References

- Bijl, H.; Schön, T. B.; van Wingerden, J.-W.; and Verhaegen, M. 2017. System identification through online sparse Gaussian process regression with input noise. *IFAC Journal of Systems and Control*, 2: 1–11.
- Bof, N.; Carli, R.; and Schenato, L. 2018. Lyapunov theory for discrete time systems. *arXiv preprint arXiv:1809.05289*.
- Chang, W.-D. 2007. Nonlinear system identification and control using a real-coded genetic algorithm. *Applied Mathematical Modelling*, 31(3): 541–550.
- Chaoui, F.-Z.; Giri*, F.; Rochdi, Y.; Haloua, M.; and Naitali, A. 2005. System identification based on Hammerstein model. *International Journal of Control*, 78(6): 430–442.
- Chen, Y.; Gu, Y.; and Hero, A. O. 2009. Sparse LMS for system identification. In *2009 IEEE international conference on acoustics, speech and signal processing*, 3125–3128. IEEE.
- Chu, H.; Wei, S.; Liu, T.; Zhao, Y.; and Miyatake, Y. 2024. Lyapunov-stable deep equilibrium models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 11615–11623.
- Finn, C.; Abbeel, P.; and Levine, S. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, 1126–1135. PMLR.
- Forgione, M.; Pura, F.; and Piga, D. 2023. From system models to class models: An in-context learning paradigm. *IEEE Control Systems Letters*, 7: 3513–3518.
- Geiselhart, R.; Gielen, R. H.; Lazar, M.; and Wirth, F. R. 2014. An alternative converse Lyapunov theorem for discrete-time systems. *Systems & Control Letters*, 70: 49–59.
- Greblicki, W.; and Pawlak, M. 2003. Identification of discrete Hammerstein systems using kernel regression estimates. *IEEE transactions on Automatic Control*, 31(1): 74–77.
- Grujić, L. T.; and Šiljak, D. 1974. Exponential stability of large-scale discrete systems. *International Journal of Control*, 19(3): 481–491.
- Hagenblad, A. 1999. Aspects of the identification of Wiener models.
- Jena, A.; Ding, F.; Wang, J.; Yao, Y.; and Xie, L. 2025. LLM-Based Adaptive Distribution Voltage Regulation Under Frequent Topology Changes: An In-Context MPC Framework. *IEEE Transactions on Smart Grid*.
- Kang, Q.; Song, Y.; Ding, Q.; and Tay, W. P. 2021. Stable neural ode with lyapunov-stable equilibrium points for defending against adversarial attacks. *Advances in Neural Information Processing Systems*, 34: 14925–14937.
- Kolter, J. Z.; and Manek, G. 2019. Learning stable deep dynamics models. *Advances in neural information processing systems*, 32.
- Korenberg, M. J. 1989. A robust orthogonal algorithm for system identification and time-series analysis. *Biological cybernetics*, 60(4): 267–276.
- Lawrence, N.; Loewen, P.; Forbes, M.; Backstrom, J.; and Gopaluni, B. 2020. Almost surely stable deep dynamics. *Advances in neural information processing systems*, 33: 18942–18953.
- Li, D.; Dong, X.; and Wang, Y. 2018. Structural system identification using sum of squares (SOS) optimization for minimizing modal dynamic residuals. *structure*, 1: 4.
- Lyashevskiy, S.; and Abel, L. A. 1994. Nonlinear systems identification using the Lyapunov method. *IFAC Proceedings Volumes*, 27(8): 271–276.
- Maciejowski, J. M. 1995. Guaranteed stability with subspace methods. *Systems & Control Letters*, 26(2): 153–156.
- Mehra, R. 1974. Optimal inputs for linear system identification. *IEEE Transactions on Automatic Control*, 19(3): 192–200.
- Ogunfunmi, T. 2007. *Adaptive nonlinear system identification: The Volterra and Wiener model approaches*. Springer Science & Business Media.
- Park, J.; Berto, F.; Jamgochian, A.; Kochenderfer, M. J.; and Park, J. 2022. Meta-sysid: A meta-learning approach for simultaneous identification and prediction. *arXiv preprint arXiv:2206.00694*.
- Pasquier, R.; and Smith, I. F. 2015. Robust system identification and model predictions in the presence of systematic uncertainty. *Advanced Engineering Informatics*, 29(4): 1096–1109.
- Pillonetto, G.; and De Nicolao, G. 2010. A new kernel-based approach for linear system identification. *Automatica*, 46(1): 81–93.
- Quaglino, A.; Gallieri, M.; Masci, J.; and Koutník, J. 2019. Snode: Spectral discretization of neural odes for system identification. *arXiv preprint arXiv:1906.07038*.
- Rockafellar, R. T.; Uryasev, S.; et al. 2000. Optimization of conditional value-at-risk. *Journal of risk*, 2: 21–42.
- Schön, T. B.; Lindsten, F.; Dahlin, J.; Wågberg, J.; Naesseth, C. A.; Svensson, A.; and Dai, L. 2015. Sequential Monte Carlo methods for system identification. *IFAC-PapersOnLine*, 48(28): 775–786.
- Stiasny, J.; Misyris, G. S.; and Chatzivasileiadis, S. 2021. Physics-informed neural networks for non-linear system identification for power system dynamics. In *2021 IEEE Madrid PowerTech*, 1–6. IEEE.
- Sun, L.; Hou, J.; Xing, C.; and Fang, Z. 2022. A robust hammerstein-wiener model identification method for highly nonlinear systems. *Processes*, 10(12): 2664.
- Tobenkin, M. M.; Manchester, I. R.; Wang, J.; Megretski, A.; and Tedrake, R. 2010. Convex optimization in identification of stable non-linear state space models. In *49th IEEE Conference on Decision and Control (CDC)*, 7232–7237. IEEE.
- Umenberger, J.; and Manchester, I. R. 2016. Scalable identification of stable positive systems. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, 4630–4635. IEEE.
- Urain, J.; Ginesi, M.; Tateo, D.; and Peters, J. 2020. Imitationflow: Learning deep stable stochastic dynamic systems by normalizing flows. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 5231–5237. IEEE.

Wu, R.-T.; and Jahanshahi, M. R. 2019. Deep convolutional neural network for structural dynamic response estimation and system identification. *Journal of Engineering Mechanics*, 145(1): 04018125.