



HAL
open science

LEADS: Learning Dynamical Systems that Generalize Across Environments

Yuan Yin, Ibrahim Ayed, Emmanuel de Bézenac, Nicolas Baskiotis, Patrick Gallinari

► **To cite this version:**

Yuan Yin, Ibrahim Ayed, Emmanuel de Bézenac, Nicolas Baskiotis, Patrick Gallinari. LEADS: Learning Dynamical Systems that Generalize Across Environments. The Thirty-Fifth Annual Conference on Neural Information Processing Systems (NeurIPS 2021), Dec 2021, Online, Unknown Region. hal-03261055v2

HAL Id: hal-03261055

<https://hal.science/hal-03261055v2>

Submitted on 16 Feb 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

LEADS: Learning Dynamical Systems that Generalize Across Environments

Yuan Yin¹, Ibrahim Ayed^{1,2}, Emmanuel de Bézenac¹, Nicolas Baskiotis¹, Patrick Gallinari^{1,3}

¹Sorbonne Université, Paris, France

²ThereSIS Lab, Thales, Paris, France ³Criteo AI Lab, Paris, France

{yuan.yin, ibrahim.ayed, emmanuel.de-bezenac, nicolas.baskiotis, patrick.gallinari}@sorbonne-universite.fr

Abstract

When modeling dynamical systems from real-world data samples, the distribution of data often changes according to the environment in which they are captured, and the dynamics of the system itself vary from one environment to another. Generalizing across environments thus challenges the conventional frameworks. The classical settings suggest either considering data as i.i.d. and learning a single model to cover all situations or learning environment-specific models. Both are sub-optimal: the former disregards the discrepancies between environments leading to biased solutions, while the latter does not exploit their potential commonalities and is prone to scarcity problems. We propose *LEADS*, a novel framework that leverages the commonalities and discrepancies among known environments to improve model generalization. This is achieved with a tailored training formulation aiming at capturing common dynamics within a shared model while additional terms capture environment-specific dynamics. We ground our approach in theory, exhibiting a decrease in sample complexity w.r.t. classical alternatives. We show how theory and practice coincides on the simplified case of linear dynamics. Moreover, we instantiate this framework for neural networks and evaluate it experimentally on representative families of nonlinear dynamics. We show that this new setting can exploit knowledge extracted from environment-dependent data and improves generalization for both known and novel environments.

1 Introduction

Data-driven approaches offer an interesting alternative and complement to physical-based methods for modeling the dynamics of complex systems and are particularly promising in a wide range of settings: e.g. if the underlying dynamics are partially known or understood, if the physical model is incomplete, inaccurate, or fails to adapt to different contexts, or if external perturbation sources and forces are not modeled. The idea of deploying machine learning (ML) to model complex dynamical systems picked momentum a few years ago, relying on recent deep learning progresses and on the development of new methods targeting the evolution of temporal and spatiotemporal systems [6, 9, 7, 21, 30, 2, 37]. It is already being applied in different scientific disciplines (see e.g. [36] for a recent survey) and could help accelerate scientific discovery to address challenging domains such as climate [32] or health [12].

However, despite promising results, current developments are limited and usually postulate an idealized setting where data is *abundant* and *the environment does not change*, the so-called “i.i.d. hypothesis”. In practice, real-world data may be expensive or difficult to acquire. Moreover, changes in the environment may be caused by many different factors. For example, in climate modeling, there are external forces (e.g. Coriolis) which depend on the spatial location [23]; or, in health science, parameters need to be personalized for each patient as for cardiac computational models [27]. More

generally, data acquisition and modeling are affected by different factors such as geographical position, sensor variability, measuring circumstances, etc. The classical paradigm either considers all the data as i.i.d. and looks for a global model, or proposes specific models for each environment. The former disregards discrepancies between the environments, thus leading to a biased solution with an averaged model which will usually perform poorly. The latter ignores the similarities between environments, thus affecting generalization performance, particularly in settings where per-environment data is limited. This is particularly problematic in dynamical settings, as small changes in initial conditions lead to trajectories not covered by the training data.

In this work, we consider a setting where it is explicitly assumed that the trajectories are collected from different environments. Note that in this setting, the i.i.d. hypothesis is removed twice: by considering the temporality of the data and by the existence of multiple environments. In many useful contexts the dynamics in each environment share similarities, while being distinct which translates into changes in the data distributions. Our objective is to leverage the similarities between environments in order to improve the modeling capacity and generalization performance, while still carefully dealing with the discrepancies across environments. This brings us to consider two research questions:

RQ1 Does modeling the differences between environments improve generalization error w.r.t. classical settings: *One-For-All*, where a unique function is trained for all environments; and *One-Per-Env.*, where a specific function is fitted for each environment? (cf. Sec. 4 for more details)

RQ2 Is it possible to extrapolate to a novel environment that has not been seen during training?

We propose LEarning Across Dynamical Systems (*LEADS*), a novel learning methodology decomposing the learned dynamics into *shared* and *environment-specific* components. The learning problem is formulated such that the *shared* component captures the dynamics common across environments and exploits all the available data, while the *environment-specific* component only models the remaining dynamics, i.e. those that cannot be expressed by the former, based on environment-specific data. We show, under mild conditions, that the learning problem is well-posed, as the resulting decomposition exists and is unique (Sec. 2.2). We then analyze the properties of this decomposition from a sample complexity perspective. While, in general, the bounds might be too loose to be practical, a more precise study is conducted in the case of linear dynamics for which theory and practice are closer. We then instantiate this framework for more general hypothesis spaces and dynamics, leading to a heuristic for the control of generalization that will be validated experimentally. Overall, we show that this framework provides better generalization properties than *One-Per-Env.*, requiring less training data to reach the same performance level (*RQ1*). The shared information is also useful to extrapolate to unknown environments: the new function for this environment can be learned from very little data (*RQ2*). We experiment with these ideas on three representative cases (Sec. 4) where the dynamics are provided by differential equations: ODEs with the Lotka-Volterra predator-prey model, and PDEs with the Gray-Scott reaction-diffusion and the more challenging incompressible Navier-Stokes equations. Experimental evidence confirms the intuition and the theoretical findings: with a similar amount of data, the approach drastically outperforms *One-For-All* and *One-Per-Env.* settings, especially in low data regimes. Up to our knowledge, it is the first time that generalization in multiple dynamical systems is addressed from an ML perspective¹.

2 Approach

2.1 Problem setting

We consider the problem of learning models of dynamical physical processes with data acquired from a set of environments E . Throughout the paper, we will assume that the dynamics in an environment $e \in E$ are defined through the evolution of differential equations. This will provide in particular a clear setup for the experiments and the validation. For a given problem, we consider that the dynamics of the different environments share common factors while each environment has its own specificity, resulting in a distinct model per environment. Both the general form of the differential equations and the specific terms of each environment are assumed to be completely unknown. x_t^e denotes the state of the equation for environment e , taking its values from a bounded set \mathcal{A} , with evolution term $f_e : \mathcal{A} \rightarrow T\mathcal{A}$, $T\mathcal{A}$ being the tangent bundle of \mathcal{A} . In other words, over a fixed time interval $[0, T]$, we have:

$$\frac{dx_t^e}{dt} = f_e(x_t^e) \tag{1}$$

¹Code is available at <https://github.com/yuan-yin/LEADS>.

We assume that, for any e , f_e lies in a functional vector space \mathcal{F} . In the experiments, we will consider one ODE, in which case $\mathcal{A} \subset \mathbb{R}^d$, and two PDEs, in which case \mathcal{A} is a d' -dimensional vector field over a bounded spatial domain $S \subset \mathbb{R}^{d'}$. The term of the data-generating dynamical system in Eq. 1 is sampled from a distribution for each e , i.e. $f_e \sim Q$. From f_e , we define \mathcal{T}_e , the data distribution of trajectories x^e verifying Eq. 1, induced by a distribution of initial states $x_0^e \sim P_0$. The data for this environment is then composed of l trajectories sampled from \mathcal{T}_e , and is denoted as $\hat{\mathcal{T}}_e$ with $x^{e,i}$ the i -th trajectory. We will denote the full dataset by $\hat{\mathcal{T}} = \bigcup_{e \in E} \hat{\mathcal{T}}_e$.

The classical empirical risk minimization (ERM) framework suggests to model the data dynamics either at the global level (*One-For-All*), taking trajectories indiscriminately from $\hat{\mathcal{T}}$, or at the specific environment level (*One-Per-Env.*), training one model for each $\hat{\mathcal{T}}_e$. Our aim is to formulate a new learning framework with the objective of explicitly considering the existence of different environments to improve the modeling strategy w.r.t. the classical ERM settings.

2.2 LEADS framework

We decompose the dynamics into two components where $f \in \mathcal{F}$ is shared across environments and $g_e \in \mathcal{F}$ is specific to the environment e , so that

$$\forall e \in E, f_e = f + g_e \quad (2)$$

Since we consider functional vector spaces, this additive hypothesis is not restrictive and such a decomposition always exists. It is also quite natural as a sum of evolution terms can be seen as the sum of the forces acting on the system. Note that the sum of two evolution terms can lead to behaviors very different from those induced by each of those terms. However, learning this decomposition from data defines an ill-posed problem: for any choice of f , there is a $\{g_e\}_{e \in E}$ such that Eq. 2 is verified. A trivial example would be $f = 0$ leading to a solution where each environment is fitted separately.

Our core idea is that f should capture as much of the shared dynamics as is possible, while g_e should focus only on the environment characteristics not captured by f . To formalize this intuition, we introduce $\Omega(g_e)$, a penalization on g_e , which precise definition will depend on the considered setting. We reformulate the learning objective as the following constrained optimization problem:

$$\min_{f, \{g_e\}_{e \in E} \in \mathcal{F}} \sum_{e \in E} \Omega(g_e) \quad \text{subject to} \quad \forall x^{e,i} \in \hat{\mathcal{T}}, \forall t, \frac{dx_t^{e,i}}{dt} = (f + g_e)(x_t^{e,i}) \quad (3)$$

Minimizing Ω aims to reduce g_e 's complexity while correctly fitting the dynamics of each environment. This argument will be made formal in the next section. Note that f will be trained on the data from all environments contrary to g_e s. A key question is then to determine under which conditions the minimum in Eq. 3 is well-defined. The following proposition provides an answer (proof cf. Sup. A):

Proposition 1 (Existence and Uniqueness). *Assume Ω is convex, then the existence of a minimal decomposition $f^*, \{g_e^*\}_{e \in E} \in \mathcal{F}$ of Eq. 3 is guaranteed. Furthermore, if Ω is strictly convex, this decomposition is unique.*

In practice, we consider the following relaxed formulation of Eq. 3:

$$\min_{f, \{g_e\}_{e \in E} \in \mathcal{F}} \sum_{e \in E} \left(\frac{1}{\lambda} \Omega(g_e) + \sum_{i=1}^l \int_0^T \left\| \frac{dx_t^{e,i}}{dt} - (f + g_e)(x_t^{e,i}) \right\|^2 dt \right) \quad (4)$$

where f, g_e are taken from a hypothesis space $\hat{\mathcal{F}}$ approximating \mathcal{F} . λ is a regularization weight and the integral term constrains the learned $f + g_e$ to follow the observed dynamics. The form of this objective and its effective calculation will be detailed in Sec. 4.4.

3 Improving generalization with LEADS

Defining an appropriate Ω is crucial for our method. In this section, we show that the generalization error should decrease with the number of environments. While the bounds might be too loose for NNs, our analysis is shown to adequately model the decreasing trend in the linear case, linking both our intuition and our theoretical analysis with empirical evidence. This then allows us to construct an appropriate Ω for NNs.

3.1 General case

After introducing preliminary notations and definitions, we define the hypothesis spaces associated with our multiple environment framework. Considering a first setting where all environments of interest are present at training time, we prove an upper-bound of their effective size based on the

covering numbers of the approximation spaces. This allows us to quantitatively control the sample complexity of our model, depending on the number of environments m and other quantities that can be considered and optimized in practice. We then consider an extension for learning on a new and unseen environment. The bounds here are inspired by ideas initially introduced in [4]. They consider multi-task classification in vector spaces, where the task specific classifiers share a common feature extractor. Our extension considers sequences corresponding to dynamical trajectories, and a model with additive components instead of function composition in their case.

Definitions. Sample complexity theory is usually defined for supervised contexts, where for a given input x we want to predict some target y . In our setting, we want to learn trajectories $(x_t^e)_{0 \leq t \leq T}$ starting from an initial condition x_0 . We reformulate this problem and cast it as a standard supervised learning problem: \mathcal{T}_e being the data distribution of trajectories for environment e , as defined in Sec. 2.1, let us consider a trajectory $x^e \sim \mathcal{T}_e$, and time $\tau \sim \text{Unif}([0, T])$; we define system states $x = x_\tau^e \in \mathcal{A}$ as input, and the corresponding values of derivatives $y = f_e(x_\tau^e) \in T\mathcal{A}$ as the associated target. We will denote \mathcal{P}_e the underlying distribution of (x, y) , and $\hat{\mathcal{P}}_e$ the associated dataset of size n .

We are searching for $f, g_e : \mathcal{A} \rightarrow T\mathcal{A}$ in an approximation function space $\hat{\mathcal{F}}$ of the original space \mathcal{F} . Let us define $\hat{\mathcal{G}} \subseteq \mathcal{F}$ the effective function space from which the g_e s are sampled. Let $f + \hat{\mathcal{G}} := \{f + g : g \in \hat{\mathcal{G}}\}$ be the hypothesis space generated by function pairs (f, g) , with a fixed $f \in \hat{\mathcal{F}}$. For any $h : \mathcal{A} \rightarrow T\mathcal{A}$, the error on some test distribution \mathcal{P}_e is given by $\text{er}_{\mathcal{P}_e}(h) = \int_{\mathcal{A} \times T\mathcal{A}} \|h(x) - y\|^2 d\mathcal{P}_e(x, y)$ and the error on the training set by $\hat{\text{er}}_{\hat{\mathcal{P}}_e}(h) = \frac{1}{n} \sum_{(x, y) \in \hat{\mathcal{P}}_e} \|h(x) - y\|^2$.

LEADS sample complexity. Let $\mathcal{C}_{\hat{\mathcal{F}}}(\varepsilon, \hat{\mathcal{F}})$ and $\mathcal{C}_{\hat{\mathcal{G}}}(\varepsilon, \hat{\mathcal{G}})$ denote the capacity of $\hat{\mathcal{F}}$ and $\hat{\mathcal{G}}$ at a certain scale $\varepsilon > 0$. Such capacity describes the approximation ability of the space. The capacity of a class of functions is defined based on covering numbers, and the precise definition is provided in Sup. B.2, Table S1. The following result is general and applies for *any* decomposition of the form $f + g_e$. It states that to guarantee a given average test error, the minimal number of samples required is a function of both capacities and the number of environments m , and it provides a step towards *RQ1* (proof see Sup. B.2):

Proposition 2. *Given m environments, let $\varepsilon_1, \varepsilon_2, \delta > 0, \varepsilon = \varepsilon_1 + \varepsilon_2$. Assume the number of examples n per environment satisfies*

$$n \geq \max \left\{ \frac{64}{\varepsilon^2} \left(\frac{1}{m} \left(\log \frac{4}{\delta} + \log \mathcal{C}_{\hat{\mathcal{G}}} \left(\frac{\varepsilon_1}{16}, \hat{\mathcal{F}} \right) \right) + \log \mathcal{C}_{\hat{\mathcal{F}}} \left(\frac{\varepsilon_2}{16}, \hat{\mathcal{G}} \right) \right), \frac{16}{\varepsilon^2} \right\} \quad (5)$$

Then with probability at least $1 - \delta$ (over the choice of training sets $\{\hat{\mathcal{P}}_e\}$), any learner $(f + g_1, \dots, f + g_m)$ will satisfy $\frac{1}{m} \sum_{e \in E} \text{er}_{\mathcal{P}_e}(f + g_e) \leq \frac{1}{m} \sum_{e \in E} \hat{\text{er}}_{\hat{\mathcal{P}}_e}(f + g_e) + \varepsilon$.

The contribution of $\hat{\mathcal{F}}$ to the sample complexity decreases as m increases, while that of $\hat{\mathcal{G}}$ remains the same: this is due to the fact that shared functions f have access to the data from all environments, which is not the case for g_e . From this finding, one infers the basis of *LEADS*: when learning from several environments, to control the generalization error through the decomposition $f_e = f + g_e$, f should account for most of the complexity of f_e while the complexity of g_e should be controlled and minimized. We then establish an explicit link to our learning problem formulation in Eq. 3. Further in this section, we will show for linear ODEs that the optimization of $\Omega(g_e)$ in Eq. 4 controls the capacity of the effective set $\hat{\mathcal{G}}$ by selecting g_e s that are as “simple” as possible.

As a corollary, we show that for a fixed total number of samples in $\hat{\mathcal{T}}$, the sample complexity will decrease as the number of environments increases. To see this, suppose that we have two situations corresponding to data generated respectively from m and m/b environments. The total sample complexity for each case will be respectively bounded by $O(\log \mathcal{C}_{\hat{\mathcal{G}}}(\frac{\varepsilon_1}{16}, \hat{\mathcal{F}}) + m \log \mathcal{C}_{\hat{\mathcal{F}}}(\frac{\varepsilon_2}{16}, \hat{\mathcal{G}}))$ and $O(b \log \mathcal{C}_{\hat{\mathcal{G}}}(\frac{\varepsilon_1}{16}, \hat{\mathcal{F}}) + m \log \mathcal{C}_{\hat{\mathcal{F}}}(\frac{\varepsilon_2}{16}, \hat{\mathcal{G}}))$. The latter being larger than the former, a situation with more environments presents a clear advantage. Fig. 4 in Sec. 4 confirms this result with empirical evidence.

LEADS sample complexity for novel environments. Suppose that problem Eq. 3 has been solved for a set of environments E , can we use the learned model for a new environment not present in the initial training set (*RQ2*)? Let e' be such a new environment, $\mathcal{P}_{e'}$ the trajectory distribution of e' , generated from dynamics $f_{e'} \sim Q$, and $\hat{\mathcal{P}}_{e'}$ an associated training set of size n' . The following results show that the number of required examples for reaching a given performance is much lower when training $f + g_{e'}$ with f fixed on this new environment than training another $f' + g_{e'}$ from scratch (proof see Sup. B.2).

Proposition 3. For all ε, δ with $0 < \varepsilon, \delta < 1$ if the number of samples n' satisfies

$$n' \geq \max \left\{ \frac{64}{\varepsilon^2} \log \frac{4\mathcal{C}(\frac{\varepsilon}{16}, f + \hat{\mathcal{G}})}{\delta}, \frac{16}{\varepsilon^2} \right\}, \quad (6)$$

then with probability at least $1 - \delta$ (over the choice of novel training set $\hat{\mathcal{P}}_{e'}$), any learner $f + g_{e'} \in f + \hat{\mathcal{G}}$ will satisfy $\text{er}_{\mathcal{P}_{e'}}(f + g_{e'}) \leq \hat{\text{er}}_{\hat{\mathcal{P}}_{e'}}(f + g_{e'}) + \varepsilon$.

In Prop. 3 as the capacity of $\hat{\mathcal{F}}$ no longer appears, the number of required samples now depends only on the capacity of $f + \hat{\mathcal{G}}$. This sample complexity is then smaller than learning from scratch $f_{e'} = f + g_{e'}$ as can be seen by comparing with Prop. 2 at $m = 1$.

From the previous propositions, it is clear that the environment-specific functions g_e need to be explicitly controlled. We now introduce a practical way to do that. Let $\omega(r, \varepsilon)$ be a strictly increasing function w.r.t. r such that

$$\log \mathcal{C}_{\hat{\mathcal{F}}}(\varepsilon, \hat{\mathcal{G}}) \leq \omega(r, \varepsilon), \quad r = \sup_{g \in \hat{\mathcal{G}}} \Omega(g) \quad (7)$$

Minimizing Ω would reduce r and then the sample complexity of our model by constraining $\hat{\mathcal{G}}$. Our goal is thus to construct such a pair (ω, Ω) . In the following, we will first show in Sec. 3.2, how one can construct a penalization term Ω based on the covering number bound for linear approximators and linear ODEs. We show with a simple use case that the generalization error obtained in practice follows the same trend as the theoretical error bound when the number of environments varies. Inspired by this result, we then propose in Sec. 3.3 an effective Ω to penalize the complexity of the neural networks g_e .

3.2 Linear case: theoretical bounds correctly predict the trend of test error

Results in Sec. 3.1 provide general guidelines for our approach. We now apply them to a linear system to see how the empirical results meet the tendency predicted by theoretical bound.

Let us consider a linear ODE $\frac{dx_t^e}{dt} = \mathbf{L}_{F_e}(x_t^e)$ where $\mathbf{L}_{F_e} : x \mapsto \mathbf{F}_e x$ is a linear transformation associated to the square real valued matrix $\mathbf{F}_e \in M_{d,d}(\mathbb{R})$. We choose as hypothesis space the space of linear functions $\hat{\mathcal{F}} \subset \mathcal{L}(\mathbb{R}^d, \mathbb{R}^d)$ and instantiate a linear LEADS $\frac{dx_t^e}{dt} = (\mathbf{L}_F + \mathbf{L}_{G_e})(x_t^e)$, $\mathbf{L}_F \in \hat{\mathcal{F}}, \mathbf{L}_{G_e} \in \hat{\mathcal{G}} \subseteq \hat{\mathcal{F}}$. As suggested in [3], we have that (proof in Sup. B.3):

Proposition 4. If for all linear maps $\mathbf{L}_{G_e} \in \hat{\mathcal{G}}, \|\mathbf{G}\|_F^2 \leq r$, if the input space is bounded s.t. $\|x\|_2 \leq b$, and the MSE loss function is bounded by c , then

$$\log \mathcal{C}_{\hat{\mathcal{F}}}(\varepsilon, \hat{\mathcal{G}}) \leq \lceil rcd(2b)^2/\varepsilon^2 \rceil \log 2d^2 =: \omega(r, \varepsilon)$$

$\omega(r, \varepsilon)$ is a strictly increasing function w.r.t. r . This indicates that we can choose $\Omega(\mathbf{L}_{G_e}) = \|\mathbf{G}\|_F$ as our optimization objective in Eq. 3. The sample complexity in Eq. 5 will decrease with the size the largest possible $r = \sup_{\mathbf{L}_{G_e} \in \hat{\mathcal{G}}} \Omega(\mathbf{L}_{G_e})$. The optimization process will reduce $\Omega(\mathbf{L}_{G_e})$ until a minimum is reached. The maximum size of the effective hypothesis space is then bounded and decreases throughout training thanks to the penalty. Then in linear case Prop. 2 becomes (proof cf. Sup. B.3):

Proposition 5. If for linear maps $\mathbf{L}_F \in \hat{\mathcal{F}}, \|\mathbf{F}\|_F^2 \leq r', \mathbf{L}_{G_e} \in \hat{\mathcal{G}}, \|\mathbf{G}\|_F^2 \leq r, \|x\|_2 \leq b$, and if the MSE loss function is bounded by c , given m environments and n samples per environment, with the probability $1 - \delta$, the generalization error upper bound is $\varepsilon = \max \{ \sqrt{(p + \sqrt{p^2 + 4q})/2}, \sqrt{16/n} \}$ where $p = \frac{64}{mn} \log \frac{4}{\delta}$ and $q = \frac{64}{n} \lceil (\frac{r'}{mz^2} + \frac{r}{(1-z)^2})cd(32b)^2 \rceil \log 2d^2$ for any $0 < z < 1$.

In Fig. 1, we take an instance of linear ODE defined by $F_e = Q\Lambda_e Q^T$ with the diagonal Λ_e specific to each environment. After solving Eq. 3 we have at the optimum that $G_e = F_e - F^* = F_e - \frac{1}{m} \sum_{e' \in E} F_{e'}$. Then we can take $r = \max_{\mathbf{L}_{G_e}} \Omega(\mathbf{L}_{G_e})$ as the norm bound of $\hat{\mathcal{G}}$ when $\Omega(g_e)$ is optimized. Fig. 1 shows on the left the test error with and without penalty and the corresponding theoretical bound on the right. We observe that, after applying the penalty Ω , the test error is reduced as well as the theoretical generalization bound, as indicated by the arrows from the dashed line to the concrete one. See Sup. B.3 for more details on this experiment.




Figure 1: Test error compared with corresponding theoretical bound. The arrows indicate the changes after applying $\Omega(g_e)$ penalty.

3.3 Nonlinear case: instantiation for neural nets

The above linear case validates the ideas introduced in Prop. 2 and provides an instantiation guide and an intuition on the more complex nonlinear case. This motivates us to instantiate the general case by choosing an appropriate approximating space $\tilde{\mathcal{F}}$ and a penalization function Ω from the generalization bounds for the corresponding space. Sup. B.4 of the Appendix contains additional details justifying those choices. For $\tilde{\mathcal{F}}$, we select the space of feed-forward neural networks with a fixed architecture. We choose the following penalty function:

$$\Omega(g_e) = \|g_e\|_\infty^2 + \alpha \|g_e\|_{\text{Lip}}^2 \quad (8)$$

where $\|g\|_\infty = \text{ess sup}|g|$ and $\|\cdot\|_{\text{Lip}}$ is the Lipschitz semi-norm, α is a hyperparameter. This is inspired by the existing capacity bound for NNs [14] (see Sup. B.4 for details). Note that constructing tight generalization bounds for neural networks is still an open research problem [26]; however, it may still yield valuable intuitions and guide algorithm design. This heuristic is tested successfully on three different datasets with different architectures in the experiments (Sec. 4).

4 Experiments

Our experiments are conducted on three families of dynamical systems described by three broad classes of differential equations. All exhibit complex and nonlinear dynamics. The first one is an ODE-driven system used for biological system modeling. The second one is a PDE-driven reaction-diffusion model, well-known in chemistry for its variety of spatiotemporal patterns. The third one is the more physically complex Navier-Stokes equation, expressing the physical laws of incompressible Newtonian fluids. To show the general validity of our framework, we will use 3 different NN architectures (MLP, ConvNet, and Fourier Neural Operator [19]). Each architecture is well-adapted to the corresponding dynamics. This also shows that the framework is valid for a variety of approximating functions.

4.1 Dynamics, environments, and datasets

Lotka-Volterra (LV). This classical model [22] is used for describing the dynamics of interaction between a predator and a prey. The dynamics follow the ODE:

$$\frac{du}{dt} = \alpha u - \beta uv, \frac{dv}{dt} = \delta uv - \gamma v$$

with u, v the number of prey and predator, $\alpha, \beta, \gamma, \delta > 0$ defining how the two species interact. The system state is $x_t^e = (u_t^e, v_t^e) \in \mathbb{R}_+^2$. The initial conditions u_0^i, v_0^i are sampled from a uniform distribution P_0 . We characterize the dynamics by $\theta = (\alpha/\beta, \gamma/\delta) \in \Theta$. An environment e is then defined by parameters θ_e sampled from a uniform distribution over a parameter set Θ . We then sample two sets of environment parameters: one used as training environments for *RQ1*, the other treated as novel environments. for *RQ2*.

Gray-Scott (GS). This reaction-diffusion model is famous for its complex spatiotemporal behavior given its simple equation formulation [29]. The governing PDE is:

$$\frac{\partial u}{\partial t} = D_u \Delta u - uv^2 + F(1 - u), \frac{\partial v}{\partial t} = D_v \Delta v + uv^2 - (F + k)v$$

where the u, v represent the concentrations of two chemical components in the spatial domain S with periodic boundary conditions, the spatially discretized state at time t is $x_t^e = (u_t^e, v_t^e) \in \mathbb{R}_+^{2 \times 32^2}$. D_u, D_v denote the diffusion coefficients respectively for u, v , and are held constant, and F, k are the reaction parameters determining the spatio-temporal patterns of the dynamics [29]. As for the initial conditions $(u_0, v_0) \sim P_0$, we consider uniform concentrations, with 3 2-by-2 squares fixed at other concentration values and positioned at uniformly sampled positions in S to trigger the reactions. An environment e is defined by its parameters $\theta_e = (F_e, k_e) \in \Theta$. We consider a set of θ_e parameters uniformly sampled from the environment distribution Q on Θ .

Navier-Stokes (NS). We consider the Navier-Stokes PDE for incompressible flows:

$$\frac{\partial w}{\partial t} = -v \cdot \nabla w + \nu \Delta w + \xi \quad \nabla \cdot v = 0$$

where v is the velocity field, $w = \nabla \times v$ is the vorticity, both v, w lie in a spatial domain S with periodic boundary conditions, ν is the viscosity and ξ is the constant forcing term in the domain S . The discretized state at time t is the vorticity $x_t^e = w_t^e \in \mathbb{R}^{32^2}$. Note that v is already contained in w . We fix $\nu = 10^{-3}$ across the environments. We sample the initial conditions $w_0^e \sim P_0$ as in [19]. An environment e is defined by its forcing term $\xi_e \in \Theta_\xi$. We uniformly sampled a set of forcing terms from Q on Θ_ξ .




Figure 2: Left: final states for GS and NS predicted by the two best baselines (*One-Per-Env.* and FT-NODE) and *LEADS* compared with ground truth. Different environment are arranged by row (3 in total). Right: the corresponding MAE error maps, the scale of the error map is $[0, 0.6]$ for GS, and $[0, 0.2]$ for NS; darker is smaller. (See Sup. D for full sequences)




Figure 3: Test predicted trajectories in phase space with two baselines (*One-Per-Env.* and FT-NODE) and *LEADS* compared with ground truth for LV for 4 envs., one per figure from left to right. Quantity of the prey u and the predator v respectively on the horizontal and the vertical axis. Initial state is the rightmost end-point of the figures and it is common to all the trajectories.

Datasets. For training, we create two datasets for LV by simulating trajectories of $K = 20$ successive points with temporal resolution $\Delta t = 0.5$. We use the first one as a set of training dynamics to validate the *LEADS* framework. We choose 10 environments and simulate 8 trajectories (thus corresponding to $n = 8 \cdot K$ data points) per environment for training. We can then easily control the number of data points and environments in experiments by taking different subsets. The second one is used to validate the improvement with *LEADS* while training on novel environments. We simulate 1 trajectory ($n = 1 \cdot K$ data points) for training. We create two datasets for further validation of *LEADS* with GS and NS. For GS, we simulate trajectories of $K = 10$ steps with $\Delta t = 40$. We choose 3 parameters and simulate 1 trajectory ($n = 1 \cdot K$ data points) for training. For NS, we simulate trajectories of $K = 10$ steps with $\Delta t = 1$. We choose 4 forcing terms and simulate 8 trajectories ($n = 8 \cdot K$ states) for training. For test-time evaluation, we create for each equation in each environment a test set of 32 trajectories ($32 \cdot K$) data points. Note that every environment dataset has the same number of trajectories and the initial conditions are fixed to equal values across the environments to ensure that the data variations only come from the dynamics themselves, i.e. for the i -th trajectory in $\hat{\mathcal{P}}_e$, $\forall e, x_0^{e,i} = x_0^i$. LV and GS data are simulated with the DOPRI5 solver in NumPy [10, 13]. NS data is simulated with the pseudo-spectral method as in [19].

4.2 Experimental settings and baselines

We validate *LEADS* in two settings: in the first one all the environments in E are available at once and then f and all the g_e s are all trained on E . In the second one, training has been performed on E as before, and we consider a novel environment $e' \notin E$: the shared term f being kept fixed, the approximating function $f_{e'} = f + g_{e'}$ is trained on the data from e' (i.e. only $g_{e'}$ is modified).

All environments available at once. We introduce five baselines used for comparing with *LEADS*: (a) *One-For-All*: learning on the entire dataset $\hat{\mathcal{P}}$ over all environments with the sum of a pair of NNs $f + g$, with the standard ERM principle, as in [2]. Although this is equivalent to use only one function f , we use this formulation to indicate that the number of parameters is the same for this experiment and for the *LEADS* ones. (b) *One-Per-Env.*: learning a specific function for each dataset $\hat{\mathcal{P}}_e$. For the same reason as above, we keep the sum formulation $(f + g)_e$. (c) Factored Tensor RNN or **FT-RNN** [33]: it modifies the recurrent neural network to integrate a one-hot environment code into each linear transformation of the network. Instead of being encoded in a separate function g_e like in *LEADS*, the environment appears here as an extra one-hot input for the RNN linear transformations. This can be implemented for representative SOTA (spatio-)temporal predictors such as GRU [8] or PredRNN [35]. (d) **FT-NODE**: a baseline for which the same environment encoding as FT-RNN is incorporated in a Neural ODE [7]. (e) Gradient-based Meta Learning or **GBML-like** method: we propose a GBML-like baseline which can directly compare to our framework. It follows the

Table 1: Results for LV, GS, and NS datasets, trained on m envs. with n data points per env.

Method	LV ($m = 10, n = 1 \cdot K$)		GS ($m = 3, n = 1 \cdot K$)		NS ($m = 4, n = 8 \cdot K$)	
	MSE train	MSE test	MSE train	MSE test	MSE train	MSE test
<i>One-For-All</i>	4.57e-1	5.08±0.56 e-1	1.55e-2	1.43±0.15 e-2	5.17e-2	7.31±5.29 e-2
<i>One-Per-Env.</i>	2.15e-5	7.95±6.96 e-3	8.48e-5	6.43±3.42 e-3	5.60e-6	1.10±0.72 e-2
FT-RNN [33]	5.29e-5	6.40±5.69 e-3	8.44e-6	8.19±3.09 e-3	7.40e-4	5.92±4.00 e-2
FT-NODE	7.74e-5	3.40±2.64 e-3	3.51e-5	3.86±3.36 e-3	1.80e-4	2.96±1.99 e-2
GBML-like	3.84e-6	5.87±5.65 e-3	1.07e-4	6.01±3.62 e-3	1.39e-4	7.37±4.80 e-3
<i>LEADS no min.</i>	3.28e-6	3.07±2.58 e-3	7.65e-5	5.53±3.43 e-3	3.20e-4	7.10±4.24 e-3
<i>LEADS</i> (Ours)	5.74e-6	1.16±0.99 e-3	5.75e-5	2.08±2.88 e-3	1.03e-4	5.95±3.65 e-3

principle of MAML [11], by training *One-For-All* at first which provides an initialization near to the given environments like GBML does, then fitting it individually for each training environment. (f) ***LEADS no min.***: ablation baseline, our proposal without the $\Omega(g_e)$ penalization. A comparison with the different baselines is proposed in Table 1 for the three dynamics. For concision, we provide a selection of results corresponding to 1 training trajectory per environment for LV and GS and 8 for NS. This is the minimal training set size for each dataset. Further experimental results when varying the number of environments from 1 to 8 are provided in Fig. 4 and Table S3 for LV.

Learning on novel environments. We consider the following training schemes with a pre-trained, fixed f : (a) ***Pre-trained-f-Only***: only the pre-trained f is used for prediction; a sanity check to ensure that f cannot predict in any novel environment without further adaptation. (b) ***One-Per-Env.***: training from scratch on $\{\mathcal{P}_{e'}\}$ as *One-Per-Env.* in the previous section. (c) ***Pre-trained-f-Plus-Trained-g_{e'}***: we train g on each dataset $\mathcal{P}_{e'}$ based on pre-trained f , i.e. $f + g_{e'}$, leaving only $g_{e'}$'s adjustable. We compare the test error evolution during training for 3 schemes above for a comparison of convergence speed and performance. Results are given in Fig. 5.

4.3 Experimental results

All environments available at once. We show the results in Table 1. For LV systems, we confirm first that the entire dataset cannot be learned properly with a single model (*One-For-All*) when the number of environments increases. Comparing with other baselines, our method *LEADS* reduces the test MSE over 85% w.r.t. *One-Per-Env.* and over 60% w.r.t. *LEADS no min.*, we also cut 50%-75% of error w.r.t. other baselines. Fig. 3 shows samples of predicted trajectories in test, *LEADS* follows very closely the ground truth trajectory, while *One-Per-Env.* under-performs in most environments. We observe the same tendency for the GS and NS systems. The error is reduced by: around 2/3 (GS) and 45% (NS) w.r.t. *One-Per-Env.*; over 60% (GS) and 15% (NS) w.r.t. *LEADS no min.*; 45-75% (GS) and 15-90% (NS) w.r.t. other baselines. In Fig. 2, the final states obtained with *LEADS* are qualitatively closer to the ground truth. Looking at the error maps on the right, we see that the errors are systematically reduced across all environments compared to the baselines. This shows that *LEADS* accumulates less errors through the integration, which suggests that *LEADS* alleviates overfitting.

We have also conducted a larger scale experiment on LV (Fig. 4) to analyze the behavior of the different training approaches as the number of environments increases. We consider three models *One-For-All*, *One-Per-Env.* and *LEADS*, 1, 2, 4 and 8 environments, and for each such case, we have 4 groups of curves, corresponding to 1, 2, 4 and 8 training trajectories per environment. We summarize the main observations. With *One-For-All* (blue), the error increases as the number of environments increases: the dynamics for each environment being indeed different, this introduces an increasingly large bias, and thus the data cannot be fit with one single model. The performance of *One-Per-Env.* (in red), for which models are trained independently for each environment, is constant as expected when the number of environments changes.




Figure 4: Test error for LV w.r.t. the number of environments. We apply the models in 1 to 8 environments. 4 groups of curves correspond to models trained with 1 to 8 trajectories per env. All groups highlight the same tendencies: increasing *One-For-All*, stable *One-Per-Env.*, and decreasing *LEADS*. More results of baselines methods in Sup. D.

LEADS (green) circumvents these issues and shows that the shared characteristics among the environments can be leveraged so as to improve generalization: it is particularly effective when the number of samples per environment is small. (See Sup. D for more details on the experiments and on the results).

Learning on novel environments. We demonstrate how the pre-trained dynamics can help to fit a model for novel environments. We took an f pre-trained by *LEADS* on a set of LV environments. Fig. 5 shows the evolution of the test loss during training for three systems: a f function pre-trained by *LEADS* on a set of LV training environments, a g_e function trained from scratch on the new environment and *LEADS* that uses a pre-trained f and learns a g_e residue on this new environment. *Pre-trained-f-Only* alone cannot predict in any novel environments. Very fast in the training stages, *Pre-trained-f-Plus-Trained- g_e* already surpasses the best error of the model trained from scratch (indicated with dotted line). Similar results are also observed with the GS and NS datasets (cf. Sup. D, Table S5). These empirical results clearly show that the learned shared dynamics accelerates and improves the learning in novel environments.




Figure 5: Test error evolution during training on 2 novel environments for LV.

4.4 Training and implementation details

Discussion on trajectory-based optimization. Solving the learning problem Eq. 2 in our setting, involves computing a trajectory loss (integral term in Eq. 4). However, in practice, we do not have access to the continuous trajectories at every instant t but only to a finite number of snapshots for the state values $\{x_{k\Delta t}\}_{0 \leq k \leq \frac{T}{\Delta t}}$ at a temporal resolution Δt . From these discrete observed trajectories, it is still possible to recover an approximate derivative $d_{k\Delta t}^\Lambda \simeq \frac{dx_{k\Delta t}}{dt}$ using a numerical scheme Λ . The integral term for a given sample in the objective Eq. 4 would then be estimated as $\sum_{k=1}^K \|d_{k\Delta t}^\Lambda - (f + g_e)(x_{\Delta tk})\|^2$. This is not the best solution and we have observed much better prediction performance for all models, including the baselines, when computing the error directly on the states, using an integral formulation $\sum_{k=1}^K \|x_{(k+1)\Delta t} - \tilde{x}_{(k+1)\Delta t}\|^2$, where $\tilde{x}_{(k+1)\Delta t}$ is the solution given by a numerical solver approximating the integral $x_{k\Delta t} + \int_{k\Delta t}^{(k+1)\Delta t} (f + g_e)(\tilde{x}_s) ds$ starting from $x_{k\Delta t}$. Comparing directly in the state space yields more accurate results for prediction as the learned network tends to correct the solver’s numerical errors, as first highlighted in [37].

Calculating Ω . Given finite data and time, the exact infinity norm and Lipschitz norm are both intractable. We opt for more practical forms in the experiments. For the infinity norm, we chose to minimize the empirical norm of the output vectors on known data points, this choice is motivated in Sup. C. In practice, we found out that dividing the output norm by its input norm works better: $\frac{1}{n} \sum_{i,k} \|g_e(x_{k\Delta t}^{e,i})\|^2 / \|x_{k\Delta t}^{e,i}\|^2$, where the $x_{k\Delta t}^{e,i}$ are known states in the training set. For the Lipschitz norm, as suggested in [5], we optimize the sum of the spectral norms of the weight at each layer $\sum_{l=1}^D \|W_l^{g_e}\|^2$. We use the power iteration method in [25] for fast spectral norm approximation.

Implementation. We used 4-layer MLPs for LV, 4-layer ConvNets for GS and Fourier Neural Operator (FNO) [19] for NS. For FT-RNN baseline, we adapted GRU [8] for LV and PredRNN [35] for GS and NS. We apply the Swish function [31] as the default activation function. Networks are integrated in time with RK4 (LV, GS) or Euler (NS), using the basic back-propagation through the internals of the solver. We apply an exponential Scheduled Sampling [17] with exponent of 0.99 to stabilize the training. We use the Adam optimizer [15] with the same learning rate 10^{-3} and $(\beta_1, \beta_2) = (0.9, 0.999)$ across the experiments. For the hyperparameters in Eq. 8, we chose respectively $\lambda = 5 \times 10^3, 10^2, 10^5$ and $\alpha = 10^{-3}, 10^{-2}, 10^{-5}$ for LV, GS and NS. All experiments are performed with a single NVIDIA Titan Xp GPU.

5 Related work

Recent approaches linking invariances to Out-of-Distribution (OoD) Generalization, such as [1, 16, 34], aim at finding a single classifier that predicts well invariantly across environments with the power of extrapolating outside the known distributions. However, in our dynamical systems context, the optimal regression function should be different in each environment, and modeling environment

bias is as important as modeling the invariant information, as both are indispensable for prediction. Thus such invariant learners are incompatible with our setting. Meta-learning methods have recently been considered for dynamical systems as in [11, 18]. Their objective is to train a single model that can be quickly adapted to a novel environment with a few data-points in limited training steps. However, in general these methods do not focus on leveraging the commonalities and discrepancies in data and may suffer from overfitting at test time [24]. Multi-task learning [38] seeks for learning shared representations of inputs that exploit the domain information. Up to our knowledge current multi-task methods have not been considered for dynamical systems. [33] apply multi-task learning for interactive physical environments but do not consider the case of dynamical systems. Other approaches like [39, 28] integrate probabilistic methods into a Neural ODE, to learn a distribution of the underlying physical processes. Their focus is on the uncertainty of a single system. [37] consider an additive decomposition but focus on the combination of physical and statistical components for a single process and not on learning from different environments.

6 Discussions

Limitations Our framework is generic and could be used in many different contexts. On the theoretical side, the existence and uniqueness properties (Prop. 1) rely on relatively mild conditions covering a large number of situations. The complexity analysis, on the other side, is only practically relevant for simple hypothesis spaces (here linear), and then serves for developing the intuition on more complex spaces (NNs here) where bounds are too loose to be informative. Another limitation is that the theory and experiments consider deterministic systems only: the experimental validation is performed on simulated deterministic data. Note however that this is the case in the vast majority of the ML literature on ODE/PDE spatio-temporal modeling [30, 20, 19, 37]. In addition, modeling complex dynamics from real world data is a problem by itself.

Conclusion We introduce *LEADS*, a data-driven framework to learn dynamics from data collected from a set of distinct dynamical systems with commonalities. Experimentally validated with three families of equations, our framework can significantly improve the test performance in every environment w.r.t. classical training, especially when the number of available trajectories is limited. We further show that the dynamics extracted by *LEADS* can boost the learning in similar new environments, which gives us a flexible framework for generalization in novel environments. More generally, we believe that this method is a promising step towards addressing the generalization problem for learning dynamical systems and has the potential to be applied to a large variety of problems.

Acknowledgements

We acknowledge financial support from the ANR AI Chairs program DL4CLIM ANR-19-CHIA-0018-01.

References

- [1] M. Arjovsky, L. Bottou, I. Gulrajani, and D. Lopez-Paz. Invariant Risk Minimization. *arXiv:1907.02893 [cs, stat]*, Mar. 2020. arXiv: 1907.02893.
- [2] I. Ayed, E. de Bézenac, A. Pajot, J. Brajard, and P. Gallinari. Learning dynamical systems from partial observations. *CoRR*, abs/1902.11136, 2019.
- [3] P. L. Bartlett, D. J. Foster, and M. J. Telgarsky. Spectrally-normalized margin bounds for neural networks. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30, pages 6240–6249. Curran Associates, Inc., 2017.
- [4] J. Baxter. A model of inductive bias learning. *J. Artif. Int. Res.*, 12(1):149–198, Mar. 2000.
- [5] A. Bietti, G. Mialon, D. Chen, and J. Mairal. A kernel perspective for regularizing deep neural networks. In K. Chaudhuri and R. Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 664–674, Long Beach, California, USA, 09–15 Jun 2019. PMLR.

- [6] S. L. Brunton, J. L. Proctor, and J. N. Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937, 2016.
- [7] R. T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud. Neural ordinary differential equations. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31, pages 6571–6583. Curran Associates, Inc., 2018.
- [8] K. Cho, B. van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, Doha, Qatar, Oct. 2014. Association for Computational Linguistics.
- [9] E. de Bézenac, A. Pajot, and P. Gallinari. Deep learning for physical processes: Incorporating prior scientific knowledge. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018.
- [10] J. Dormand and P. Prince. A family of embedded runge-kutta formulae. *Journal of Computational and Applied Mathematics*, 6(1):19 – 26, 1980.
- [11] C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. *CoRR*, abs/1703.03400, 2017.
- [12] S. Fresca, A. Manzoni, L. Dedè, and A. Quarteroni. *Deep learning-based reduced order models in cardiac electrophysiology*, volume 15. 2020.
- [13] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. F. del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, and T. E. Oliphant. Array programming with NumPy. *Nature*, 585(7825):357–362, Sept. 2020.
- [14] D. Haussler. Decision theoretic generalizations of the pac model for neural net and other learning applications. *Information and Computation*, 100(1):78 – 150, 1992.
- [15] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In Y. Bengio and Y. LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [16] D. Krueger, E. Caballero, J. Jacobsen, A. Zhang, J. Binas, R. Le Priol, and A. C. Courville. Out-of-distribution generalization via risk extrapolation (rex). *CoRR*, abs/2003.00688, 2020.
- [17] A. Lamb, A. Goyal, Y. Zhang, S. Zhang, A. Courville, and Y. Bengio. Professor Forcing: A New Algorithm for Training Recurrent Networks. *arXiv:1610.09038 [cs, stat]*, Oct. 2016. arXiv: 1610.09038.
- [18] S. Lee, H. Yang, and W. Seong. Identifying physical law of hamiltonian systems via meta-learning. *CoRR*, abs/2102.11544, 2021.
- [19] Z. Li, N. B. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart, and A. Anandkumar. Fourier neural operator for parametric partial differential equations. In *International Conference on Learning Representations*, 2021.
- [20] Z. Long, Y. Lu, and B. Dong. Pde-net 2.0: Learning pdes from data with A numeric-symbolic hybrid deep network. *CoRR*, abs/1812.04426, 2018.
- [21] Z. Long, Y. Lu, X. Ma, and B. Dong. Pde-net: Learning pdes from data. In *International Conference on Machine Learning*, pages 3214–3222, 2018.
- [22] A. J. Lotka. Elements of physical biology. *Science Progress in the Twentieth Century (1919-1933)*, 21(82):341–343, 1926.

- [23] G. Madec, R. Bourdallé-Badie, J. Chanut, E. Clementi, A. Coward, C. Ethé, D. Iovino, D. Lea, C. Lévy, T. Lovato, N. Martin, S. Masson, S. Mocavero, C. Rousset, D. Storkey, M. Vancoppenolle, S. Müeller, G. Nurser, M. Bell, and G. Samson. Nemo ocean engine, Oct. 2019. Add SI3 and TOP reference manuals.
- [24] N. Mishra, M. Rohaninejad, X. Chen, and P. Abbeel. Meta-learning with temporal convolutions. *CoRR*, abs/1707.03141, 2017.
- [25] T. Miyato, T. Kataoka, M. Koyama, and Y. Yoshida. Spectral normalization for generative adversarial networks. *CoRR*, abs/1802.05957, 2018.
- [26] V. Nagarajan and J. Z. Kolter. Uniform convergence may be unable to explain generalization in deep learning. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [27] A. Neic, F. O. Campos, A. J. Prassl, S. A. Niederer, M. J. Bishop, E. J. Vigmond, and G. Plank. Efficient computation of electrograms and egs in human whole heart simulations using a reaction-eikonal model. *Journal of Computational Physics*, 346:191 – 211, 2017.
- [28] A. Norcliffe, C. Bodnar, B. Day, J. Moss, and P. Liò. Neural ODE processes. In *International Conference on Learning Representations*, 2021.
- [29] J. E. Pearson. Complex patterns in a simple system. *Science*, 261(5118):189–192, 1993.
- [30] M. Raissi, P. Perdikaris, and G. E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
- [31] P. Ramachandran, B. Zoph, and Q. V. Le. Searching for activation functions. *CoRR*, abs/1710.05941, 2017.
- [32] M. Reichstein, G. Camps-Valls, B. Stevens, M. Jung, J. Denzler, N. Carvalhais, and Prabhat. Deep learning and process understanding for data-driven Earth system science. *Nature*, 566:195–204, 2019.
- [33] S. Spieckermann, S. Düll, S. Udluft, A. Hentschel, and T. Runkler. Exploiting similarity in system identification tasks with recurrent neural networks. *Neurocomputing*, 169:343 – 349, 2015. Learning for Visual Semantic Understanding in Big Data ESANN 2014 Industrial Data Processing and Analysis.
- [34] D. Teney, E. Abbasnejad, and A. van den Hengel. Unshuffling Data for Improved Generalization. 2020.
- [35] Y. Wang, M. Long, J. Wang, Z. Gao, and P. S. Yu. Predrnn: Recurrent neural networks for predictive learning using spatiotemporal lstms. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [36] J. D. Willard, X. Jia, S. Xu, M. Steinbach, and V. Kumar. Integrating physics-based modeling with machine learning: A survey. volume 1, pages 1–34, 2020.
- [37] Y. Yin, V. Le Guen, J. Dona, E. de Bezenac, I. Ayed, N. Thome, and P. Gallinari. Augmenting physical models with deep networks for complex dynamics forecasting. In *International Conference on Learning Representations*, 2021.
- [38] Y. Zhang and Q. Yang. A survey on multi-task learning. *CoRR*, abs/1707.08114, 2017.
- [39] Çağatay Yıldız, M. Heinonen, and H. Lähdesmäki. Ode²vae: Deep generative second order odes with bayesian neural networks, 2019.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes] In Sec. 6.
 - (c) Did you discuss any potential negative societal impacts of your work? [No]
The only relevant societal impact is around the computational cost, while we use very limited computation power (maximum with a single GPU).
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes]
In Sec. 2.1 and Sec. 3.
 - (b) Did you include complete proofs of all theoretical results? [Yes]
In the supplemental material.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
We will provide the code in the supplemental material.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] In Sec. 4.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] In Sec. 4.
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] In Sec. 4.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [Yes] In Sec. 4.
 - (b) Did you mention the license of the assets? [Yes] In the supplemental material.
 - (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

LEADS: Learning Dynamical Systems that Generalize Across Environments

Supplemental Material

Yuan Yin¹, Ibrahim Ayed^{1,2}, Emmanuel de Bézenac¹, Nicolas Baskiotis¹, Patrick Gallinari^{1,3}

¹Sorbonne Université, Paris, France

²ThereSIS Lab, Thales, Paris, France ³Criteo AI Lab, Paris, France

{yuan.yin, ibrahim.ayed, emmanuel.de-bezenac,
nicolas.baskiotis,patrick.gallinari}@sorbonne-universite.fr

A Proof of Proposition 1

Proposition 1 (Existence and Uniqueness). *Assume Ω is convex, then the existence of a minimal decomposition $f^*, \{g_e^*\}_{e \in E} \in \mathcal{F}$ of Eq. 3 is guaranteed. Furthermore, if Ω is strictly convex, this decomposition is unique.*

Proof. The optimization problem is:

$$\min_{f, g_e \in \mathcal{F}} \sum_{e \in E} \Omega(g_e) \quad \text{subject to} \quad \forall x^{e,i} \in \hat{\mathcal{T}}, \forall t, \frac{dx_t^{e,i}}{dt} = (f + g_e)(x_t^{e,i}) \quad (3)$$

The idea is to first reconstruct the full functional from the trajectories of $\hat{\mathcal{T}}$. By definition, \mathcal{A}^e is the set of points reached by trajectories in $\hat{\mathcal{T}}$ from environment e so that:

$$\mathcal{A}^e = \{x \in \mathbb{R}^d \mid \exists x_t^e \in \hat{\mathcal{T}}, \exists t, x_t^e = x\}$$

Then let us define a function f_e^{data} in the following way, $\forall e \in E$, take $a \in \mathcal{A}^e$, we can find $x^e \in \hat{\mathcal{T}}$ and t_0 such that $x_{t_0}^e = a$. Differentiating x^e at t_0 , which is possible by definition of $\hat{\mathcal{T}}$, we take:

$$f_e^{\text{data}}(a) = \left. \frac{dx_t^e}{dt} \right|_{t=t_0}$$

For any (f, g_e) satisfying the constraint in Eq. 3, we then have $(f + g_e)(a) = \left. \frac{dx_t}{dt} \right|_{t_0} = f_e^{\text{data}}(a)$ for all $a \in \mathcal{A}^e$. Conversely, any pair such that $(f, g_e) \in \mathcal{F} \times \mathcal{F}$ and $f + g_e = f_e^{\text{data}}$, verifies the constraint.

Thus we have the equivalence between Eq. 3 and the following objective:

$$\min_{f \in \mathcal{F}} \sum_e \Omega(f_e^{\text{data}} - f) \quad (S1)$$

The result directly follows from the fact that the objective is a sum of (strictly) convex functions in f and is thus (strictly) convex in f . \square

B Further details on the generalization with LEADS

In this section, we will give more details on the link between our framework and its generalization performance. After introducing the necessary definitions in Sec. B.1, we show the proofs of the results for the general case in Sec. 3. Then in Sec. B.3 we provide the instantiation for linear approximators. Finally, we show how we derived our heuristic instantiation for neural networks in Eq. 8 in Sec. 3.3 from the existing capacity bound for neural networks.

Table S1: Capacity definitions of different sets by covering number with associated metric or pseudo-metric.

Capacity	Metric or pseudo-metric	Mentioned in
$\mathcal{C}(\varepsilon, \mathbb{H}^m) := \sup_{\mathcal{P}} \mathcal{N}(\varepsilon, \mathbb{H}^m, d_{\mathcal{P}})$	$d_{\mathcal{P}}((f + g_1, \dots, f + g_m), (f' + g'_1, \dots, f' + g'_m)) = \int_{(\mathcal{A} \times T \mathcal{A})^m} \frac{1}{m} \sum_{e \in E} \ (f + g_e)(x^e) - y^e\ ^2 - \sum_{e \in E} \ (f' + g'_e)(x^e) - y^e\ ^2 d\mathcal{P}(\mathbf{x}, \mathbf{y})$	Theorem S1; Prop. S1
$\mathcal{C}_{\hat{\mathcal{F}}}(\varepsilon, \hat{\mathcal{F}}) := \sup_{\mathcal{P}} \mathcal{N}(\varepsilon, \hat{\mathcal{F}}, d_{[\mathcal{P}, \hat{\mathcal{F}}]})$	$d_{[\mathcal{P}, \hat{\mathcal{F}}]}(f, f') = \int_{\mathcal{A} \times T \mathcal{A}} \sup_{g \in \hat{\mathcal{G}}} \ (f + g)(x) - y\ ^2 - \ (f' + g)(x) - y\ ^2 d\mathcal{P}(x, y)$	Prop. 2, S1, S3; Cor. S1
$\mathcal{C}_{\hat{\mathcal{F}}}(\varepsilon, \hat{\mathcal{G}}) := \sup_{\mathcal{P}} \mathcal{N}(\varepsilon, \hat{\mathcal{G}}, d_{[\mathcal{P}, \hat{\mathcal{F}}]})$	$d_{[\mathcal{P}, \hat{\mathcal{F}}]}(g, g') = \int_{\mathcal{A} \times T \mathcal{A}} \sup_{f \in \hat{\mathcal{F}}} \ (f + g)(x) - y\ ^2 - \ (f + g')(x) - y\ ^2 d\mathcal{P}(x, y)$	Prop. 2, S1, S2
$\mathcal{C}(\varepsilon, f + \hat{\mathcal{G}}) := \sup_{\mathcal{P}} \mathcal{N}(\varepsilon, f + \hat{\mathcal{G}}, d_{\mathcal{P}})$	$d_{\mathcal{P}}(f + g, f + g') = \int_{\mathcal{A} \times T \mathcal{A}} \ (f + g)(x) - y\ ^2 - \ (f + g')(x) - y\ ^2 d\mathcal{P}(x, y)$	Prop. 3
$\mathcal{C}(\varepsilon, \hat{\mathcal{G}}, L^1) := \sup_{\mathcal{P}} \mathcal{N}(\varepsilon, \hat{\mathcal{G}}, d_{L^1(\mathcal{P})})$	$d_{L^1(\mathcal{P})}(g, g') = \int_{\mathbb{R}^d} \ (g - g')(x)\ _1 d\mathcal{P}(x)$	Prop. S2; Theorem S3
$\mathcal{C}(\varepsilon, \hat{\mathcal{G}}, L^2) := \sup_{\mathcal{P}} \mathcal{N}(\varepsilon, \hat{\mathcal{G}}, d_{L^2(\mathcal{P})})$	$d_{L^2(\mathcal{P})}(g, g') = \sqrt{\int_{\mathbb{R}^d} \ (g - g')(x)\ _2^2 d\mathcal{P}(x)}$	Prop. 4; Lemma S1

B.1 Preliminaries

Table S1 gives the definition of the different capacity instances considered in the paper for each hypothesis space, and the associated distances. We say that a space \mathcal{H} is ε -covered by a set H , with respect to a metric or pseudo-metric $d(\cdot, \cdot)$, if for all $h \in \mathcal{H}$ there exists $h' \in H$ with $d(h, h') \leq \varepsilon$. We define by $\mathcal{N}(\varepsilon, \mathcal{H}, d)$ the cardinality of the smallest H that ε -covers \mathcal{H} , also called covering number [S7]. The capacity of each hypothesis space is then defined by the maximum covering number over all distributions. Note that the loss function is involved in every metric in Table S1. For simplicity, we therefore omit the notation of loss function for the hypothesis spaces.

As in [S2], covering numbers are based on pseudo-metrics. We can verify that all distances in Table S1 are pseudo-metrics:

Proof. This is trivially verified. For example, for the distance $d_{\mathcal{P}}(f + g, f + g')$ given in Table S1, which is the distance between $f + g, f + g' \in f + \hat{\mathcal{G}}$, it is easy to check that the following properties do hold:

- $d_{\mathcal{P}}(f + g, f + g') = 0$ (subtraction of same functions evaluated on same x and y)
- $d_{\mathcal{P}}(f + g, f + g') = d_{\mathcal{P}}(f + g', f + g)$ (evenness of absolute value)
- $d_{\mathcal{P}}(f + g, f + g') \leq d_{\mathcal{P}}(f + g, f + g'') + d_{\mathcal{P}}(f + g'', f + g')$ (triangular inequality of absolute value)

Other distances in Table S1 can be proven to be pseudo-metrics in the same way. \square

B.2 General Case

B.2.1 Proof of Proposition 2

Proposition 2. *Given m environments, let $\varepsilon_1, \varepsilon_2, \delta > 0, \varepsilon = \varepsilon_1 + \varepsilon_2$. Assume the number of examples n per environment satisfies*

$$n \geq \max \left\{ \frac{64}{\varepsilon^2} \left(\frac{1}{m} \left(\log \frac{4}{\delta} + \log \mathcal{C}_{\hat{\mathcal{F}}} \left(\frac{\varepsilon_1}{16}, \hat{\mathcal{F}} \right) \right) + \log \mathcal{C}_{\hat{\mathcal{F}}} \left(\frac{\varepsilon_2}{16}, \hat{\mathcal{G}} \right) \right), \frac{16}{\varepsilon^2} \right\} \quad (5)$$

Then with probability at least $1 - \delta$ (over the choice of training sets $\{\hat{\mathcal{P}}_e\}$), any learner $(f + g_1, \dots, f + g_m)$ will satisfy $\frac{1}{m} \sum_{e \in E} \text{er}_{\mathcal{P}_e}(f + g_e) \leq \frac{1}{m} \sum_{e \in E} \hat{\text{er}}_{\hat{\mathcal{P}}_e}(f + g_e) + \varepsilon$.

Proof. We introduce some extra definitions that are necessary for proving the proposition. Let $\mathcal{H} = f + \hat{\mathcal{G}}$ defined for each $f \in \hat{\mathcal{F}}$, and let us define the product space $\mathcal{H}^m = \{(f + g_1, \dots, f + g_m) : f + g_e \in \mathcal{H}\}$. Functions in this hypothesis space all have the same f , but not necessarily the same g_e . Let \mathbb{H} be the collection of all hypothesis spaces $\mathcal{H} = f + \hat{\mathcal{G}}, \forall f \in \hat{\mathcal{F}}$. The hypothesis space associated to multiple environments is then defined as $\mathbb{H}^m := \bigcup_{\mathcal{H} \in \mathbb{H}} \mathcal{H}^m$.

Our proof makes use of two intermediary results addressed in Theorem S1 and Prop. S1.

Theorem S1 ([S2], Theorem 4, adapted to our setting). *Assuming \mathbb{H} is a permissible hypothesis space family. For all $\varepsilon > 0$, if the number of examples n of each environment satisfies:*

$$n \geq \max \left\{ \frac{64}{m\varepsilon^2} \log \frac{4\mathcal{C}(\frac{\varepsilon}{16}, \mathbb{H}^m)}{\delta}, \frac{16}{\varepsilon^2} \right\}$$

Then with probability at least $1 - \delta$ (over the choice of $\{\hat{\mathcal{P}}_e\}$), any $(f + g_1, \dots, f + g_m)$ will satisfy

$$\frac{1}{m} \sum_{e \in E} \text{er}_{\mathcal{P}_e}(f + g_e) \leq \frac{1}{m} \sum_{e \in E} \hat{\text{er}}_{\hat{\mathcal{P}}_e}(f + g_e) + \varepsilon$$

Note that permissibility (as defined in [S2]) is a weak measure-theoretic condition satisfied by many real world hypothesis space families [S2]. We will now express the capacity of \mathbb{H}^m in terms of the capacities of its two constituent component-spaces $\hat{\mathcal{F}}$ and $\hat{\mathcal{G}}$, thus leading to the main result.

Proposition S1. *For all $\varepsilon, \varepsilon_1, \varepsilon_2 > 0$ such that $\varepsilon = \varepsilon_1 + \varepsilon_2$,*

$$\log \mathcal{C}(\varepsilon, \mathbb{H}^m) \leq \log \mathcal{C}_{\hat{\mathcal{G}}}(\varepsilon_1, \hat{\mathcal{F}}) + m \log \mathcal{C}_{\hat{\mathcal{F}}}(\varepsilon_2, \hat{\mathcal{G}}) \quad (\text{S2})$$

Proof of Proposition S1. To prove the proposition it is sufficient to show the property of covering sets for any joint distribution defined on all environments \mathcal{P} on the space $(\mathcal{A} \times T\mathcal{A})^m$. Let us then fix such a distribution \mathcal{P} . and let $\bar{\mathcal{P}} = \frac{1}{m} \sum_{e \in E} \mathcal{P}_e$ be the average distribution.

Suppose that F is an ε_1 -cover of $(\hat{\mathcal{F}}, d_{[\bar{\mathcal{P}}, \hat{\mathcal{G}}]})$ and $\{G_e\}_{e \in E}$ are ε_2 -covers of $(\hat{\mathcal{G}}, d_{[\mathcal{P}_e, \hat{\mathcal{F}}]})$. Let $H = \{(x_1, \dots, x_m) \mapsto ((f + g_1)(x_1), \dots, (f + g_m)(x_m)), f \in F, g_e \in G_e\}$, be a set built from the covering sets aforementioned. Note that by definition $|H| = |F| \cdot \prod_{e \in E} |G_e| \leq \mathcal{C}_{\hat{\mathcal{G}}}(\varepsilon_1, \hat{\mathcal{F}}) \mathcal{C}_{\hat{\mathcal{F}}}(\varepsilon_2, \hat{\mathcal{G}})^m$ as we take some distribution instances.

For each learner $(f + g_1, \dots, f + g_m) \in \mathbb{H}^m$ in the hypothesis space, we take any $f' \in F$ such that $d_{[\bar{\mathcal{P}}, \hat{\mathcal{G}}]}(f, f') \leq \varepsilon_1$ and $g'_e \in G_e$ for all e such that $d_{[\mathcal{P}_e, \hat{\mathcal{F}}]}(g_e, g'_e) \leq \varepsilon_2$, and we build $(f' + g'_1, \dots, f' + g'_m)$. The distance is then:

$$\begin{aligned} & d_{\mathcal{P}}((f + g_1, \dots, f + g_m), (f' + g'_1, \dots, f' + g'_m)) \\ & \leq d_{\mathcal{P}}((f + g_1, \dots, f + g_m), (f' + g_1, \dots, f' + g_m)) \\ & \quad + d_{\mathcal{P}}((f' + g_1, \dots, f' + g_m), (f' + g'_1, \dots, f' + g'_m)) \\ & \hspace{15em} (\text{triangular inequality of pseudo-metric}) \\ & \leq \frac{1}{m} \left[\sum_{e \in E} d_{\mathcal{P}_e}(f + g_e, f' + g_e) + \sum_{e \in E} d_{\mathcal{P}_e}(f' + g_e, f' + g'_e) \right] \\ & \hspace{15em} (\text{triangular inequality of absolute value}) \\ & \leq \frac{1}{m} \sum_{e \in E} d_{[\mathcal{P}_e, \hat{\mathcal{G}}]}(f, f') + \frac{1}{m} \sum_{e \in E} d_{[\mathcal{P}_e, \hat{\mathcal{F}}]}(g_e, g'_e) \quad (\text{by definition of } d_{[\mathcal{P}_e, \hat{\mathcal{G}}]} \text{ and } d_{[\mathcal{P}_e, \hat{\mathcal{F}}]}) \\ & = d_{[\bar{\mathcal{P}}, \hat{\mathcal{G}}]}(f, f') + \frac{1}{m} \sum_{e \in E} d_{[\mathcal{P}_e, \hat{\mathcal{F}}]}(g_e, g'_e) \leq \varepsilon_1 + \varepsilon_2 \\ & \hspace{15em} (\text{mean of the distance on different } \mathcal{P}_e \text{ is the distance on } \bar{\mathcal{P}}) \end{aligned}$$

To conclude, for any distribution \mathcal{P} , when F is an ε_1 -cover of $\hat{\mathcal{F}}$ and $\{G_e\}$ are ε_2 -covers of $\hat{\mathcal{G}}$, the set H built upon them is an $(\varepsilon_1 + \varepsilon_2)$ -cover of \mathbb{H}^m . Then if we take the maximum over all distributions we conclude that $\mathcal{C}(\varepsilon_1 + \varepsilon_2, \mathbb{H}^m) \leq \mathcal{C}_{\hat{\mathcal{G}}}(\varepsilon_1, \hat{\mathcal{F}}) \mathcal{C}_{\hat{\mathcal{F}}}(\varepsilon_2, \hat{\mathcal{G}})^m$ and we have Eq. S2. ■

We can now use the bound developed in Prop. S1 and use it together with Theorem S1, therefore concluding the proof of Prop. 2. \square

B.2.2 Proof of Proposition 3

Proposition 3. For all ε, δ with $0 < \varepsilon, \delta < 1$ if the number of samples n' satisfies

$$n' \geq \max \left\{ \frac{64}{\varepsilon^2} \log \frac{4\mathcal{C}(\frac{\varepsilon}{16}, f + \hat{\mathcal{G}}), 16}{\delta}, \frac{16}{\varepsilon^2} \right\}, \quad (6)$$

then with probability at least $1 - \delta$ (over the choice of novel training set $\hat{\mathcal{P}}_{e'}$), any learner $f + g_{e'} \in f + \hat{\mathcal{G}}$ will satisfy $\text{er}_{\mathcal{P}_{e'}}(f + g_{e'}) \leq \hat{\text{er}}_{\hat{\mathcal{P}}_{e'}}(f + g_{e'}) + \varepsilon$.

Proof. The proof is derived from the following theorem which can be easily adapted to our context:

Theorem S2 ([S2], Theorem 3). Let \mathcal{H} a permissible hypothesis space. For all $0 < \varepsilon, \delta < 1$, if the number of examples n of each environment satisfies:

$$n \geq \max \left\{ \frac{64}{m\varepsilon^2} \log \frac{4\mathcal{C}(\frac{\varepsilon}{16}, \mathcal{H})}{\delta}, \frac{16}{\varepsilon^2} \right\}$$

Then with probability at least $1 - \delta$ (over the choice of dataset $\hat{\mathcal{P}}$ sampled from \mathcal{P}), any $h \in \mathcal{H}$ will satisfy

$$\text{er}_{\mathcal{P}}(h) \leq \hat{\text{er}}_{\hat{\mathcal{P}}}(h) + \varepsilon$$

Given that $\hat{\mathcal{P}}_{e'}$ is sampled from the same environment distribution Q , then by fixing the pre-trained f , we fix the space of hypothesis to $f + \hat{\mathcal{G}}$, and we apply the Theorem S2 to obtain the proposition. \square

B.3 Linear case

We provide here the proofs of theoretical bounds given in Sec. 3.2. See the description in Sup. D for the detailed information on the example linear ODE dataset and the training with varying number of environments.

B.3.1 Proof of Proposition 4

Proposition 4. If for all linear maps $L_{G_e} \in \hat{\mathcal{G}}$, $\|G\|_F^2 \leq r$, if the input space is bounded s.t. $\|x\|_2 \leq b$, and the MSE loss function is bounded by c , then

$$\log \mathcal{C}_{\hat{\mathcal{G}}}(\varepsilon, \hat{\mathcal{G}}) \leq \lceil rcd(2b)^2/\varepsilon^2 \rceil \log 2d^2 =: \omega(r, \varepsilon)$$

Proof. Let us take G an $\frac{\varepsilon}{2\sqrt{c}}$ -cover of $\hat{\mathcal{G}}$ with L^2 -distance: $d_{L^2(\mathcal{P})}$ (see definition in Table S1).

Therefore, for each $L_G \in \hat{\mathcal{G}}$ take $g' \in G$ such that $d_{L^2}(L_G, L_{G'}) \leq \frac{\varepsilon}{2\sqrt{c}}$, then

$$\begin{aligned} & d_{[\mathcal{P}, \hat{\mathcal{G}}]}(L_G, L_{G'}) \\ &= \int_{\mathcal{A} \times \mathcal{A}'} \sup_{L_F \in \hat{\mathcal{F}}} \left| \|(F + G)x - y\|_2^2 - \|(F + G')x - y\|_2^2 \right| d\mathcal{P}(x, y) \\ &\leq \int_{\mathcal{A} \times T\mathcal{A}} \sup_{L_F \in \hat{\mathcal{F}}} \|(G - G')x\|_2 (\|(F + G)x - y\|_2 + \|(F + G')(x) - y\|_2) d\mathcal{P}(x, y) \\ &\leq \sqrt{\int_{\mathcal{A}} \|(G - G')x\|_2^2 d\mathcal{P}(x)} \sqrt{\int_{\mathcal{A} \times T\mathcal{A}} \sup_{L_F \in \hat{\mathcal{F}}} (\|(F + G)x - y\|_2 + \|(F + G')(x) - y\|_2)^2 d\mathcal{P}(x, y)} \\ &\leq 2\sqrt{c} \sqrt{\int_{\mathbb{R}^d} \|(G - G')x\|_2^2 d\mathcal{P}(x)} \leq \varepsilon \end{aligned}$$

We have the $\mathcal{C}_{\mathcal{F}}(\varepsilon, \hat{\mathcal{G}}) \leq \mathcal{C}(\frac{\varepsilon}{2\sqrt{c}}, \hat{\mathcal{G}}, L^2)$. According to the following lemma:

Lemma S1 ([S1], Lemma 3.2, Adapted). Given positive reals (a, b, ε) and positive integer d . Let vector $x \in \mathbb{R}^d$ be given with $\|x\|_p \leq b$, $\hat{\mathcal{G}} = \{L_G : G \in \mathbb{R}^{d \times d}, \|G\|_F^2 \leq r\}$ where $\|\cdot\|_F$ is the

Frobenius norm. Then

$$\log \mathcal{C}(\varepsilon, \hat{\mathcal{G}}, L^2) \leq \left\lceil \frac{rdb^2}{\varepsilon^2} \right\rceil \log 2d^2$$

And we obtain that

$$\log \mathcal{C}_{\hat{\mathcal{F}}}(\varepsilon, \hat{\mathcal{G}}) \leq \left\lceil \frac{rcd(2b)^2}{\varepsilon^2} \right\rceil \log 2d^2 =: \omega(r, \varepsilon)$$

where $\omega(r, \varepsilon)$ is a strictly increasing function w.r.t. r . □

B.3.2 Proof of Proposition 5

Proposition 5. *If for linear maps $L_F \in \hat{\mathcal{F}}$, $\|F\|_F^2 \leq r'$, $L_G \in \hat{\mathcal{G}}$, $\|G\|_F^2 \leq r$, $\|x\|_2 \leq b$, and if the MSE loss function is bounded by c , given m environments and n samples per environment, with the probability $1 - \delta$, the generalization error upper bound is $\varepsilon = \max \{ \sqrt{(p + \sqrt{p^2 + 4q})/2}, \sqrt{16/n} \}$ where $p = \frac{64}{mn} \log \frac{4}{\delta}$ and $q = \frac{64}{n} \lceil (\frac{r'}{mz^2} + \frac{r}{(1-z)^2}) cd(32b)^2 \rceil \log 2d^2$ for any $0 < z < 1$.*

Proof. This can be derived from Prop. 2 with the help of Prop. 4 for linear maps. If we take the lower bounds of two capacities $\log \mathcal{C}_{\hat{\mathcal{F}}}(\frac{\varepsilon_1}{16}, \hat{\mathcal{G}})$ and $\log \mathcal{C}_{\hat{\mathcal{G}}}(\frac{\varepsilon_2}{16}, \hat{\mathcal{F}})$ for the linear maps hypothesis spaces $\hat{\mathcal{F}}, \hat{\mathcal{G}}$, then the number of required samples per environment n now can be expressed as follows:

$$n = \max \left\{ \frac{64}{\varepsilon^2} \left(\frac{1}{m} \log \frac{4}{\delta} + \frac{1}{m} \left\lceil \frac{r'cd(32b)^2}{\varepsilon_1^2} \right\rceil \log 2d^2 + \left\lceil \frac{rcd(32b)^2}{\varepsilon_2^2} \right\rceil \log 2d^2 \right), \frac{16}{\varepsilon^2} \right\}$$

To simplify the resolution of the equation above, we take $\varepsilon_1 = z\varepsilon$ for any $0 < z < 1$, then $\varepsilon_2 = \varepsilon - \varepsilon_1 = (1 - z)\varepsilon$. Then by resolving the equation, the generalization margin is then upper bounded by ε with:

$$\varepsilon = \max \left\{ \sqrt{\frac{p + \sqrt{p^2 + 4q}}{2}}, \sqrt{\frac{16}{n}} \right\}$$

where $p = \frac{64}{mn} \log \frac{4}{\delta}$ and $q = \frac{64}{n} \lceil (\frac{r}{mz^2} + \frac{r'}{(1-z)^2}) cd(32b)^2 \rceil \log 2d^2$. □

B.4 Nonlinear case: instantiation for neural networks

We show in this section how we design a concrete model for nonlinear dynamics following the general guidelines given in Sec. 3.1. This is mainly composed of the following two parts: (a) choosing an appropriate approximation space and (b) choosing a penalization function Ω for this space. It is important to note that, even if the bounds given in the following sections may be loose in general, it could provide useful intuitions on the design of the algorithms which can be validated by experiments in our case.

B.4.1 Choosing approximation space $\hat{\mathcal{F}}$

We choose the space of feed-forward neural networks with a fixed architecture. Given the universal approximation properties of neural networks [S5], and the existence of efficient optimization algorithms [S3], this is a reasonable choice, but other families of approximating functions could be used as well.

We then consider the function space of neural networks with D -layers with inputs and outputs in \mathbb{R}^d : $\hat{\mathcal{F}}_{\text{NN}} = \{ \nu : x \mapsto \sigma_D(W_D \cdots \sigma_1(W_1 x)) : x, \nu(x) \in \mathbb{R}^d \}$, D is the depth of the network, σ_j is a Lipschitz activation function at layer j , and W_j weight matrix from layer $j - 1$ to j . The number of adjustable parameters is fixed to W for the architecture. This definition covers fully connected NNs and convolutional NNs. Note that the Fourier Neural Operator [S6] used in the experiments for NS can be also covered by the definition above, as it performs alternatively the convolution in the Fourier space.

B.4.2 Choosing penalization Ω

Now we choose an Ω for the space above. Let us first introduce a practical way to bound the capacity of $\hat{\mathcal{G}} \in \hat{\mathcal{F}}_{\text{NN}}$. Proposition S2 tells us that for a fixed NN architecture (implying constant parameter

number W and depth D), we can control the capacity through the maximum output norm R and Lipschitz norm L defined in the proposition.

Proposition S2. *If for all neural network $g \in \hat{\mathcal{G}}$, $\|g\|_\infty = \text{ess sup}|g| \leq R$ and $\|g\|_{\text{Lip}} \leq L$, with $\|\cdot\|_{\text{Lip}}$ the Lipschitz semi-norm, then:*

$$\log \mathcal{C}_{\hat{\mathcal{F}}}(\varepsilon, \hat{\mathcal{G}}) \leq \omega(R, L, \varepsilon) \quad (\text{S3})$$

where $\omega(R, L, \varepsilon) = c_1 \log \frac{RL}{\varepsilon} + c_2$ for $c_1 = 2W$ and $c_2 = 2W \log 8e\sqrt{c}D$, with c the bound of MSE loss. $\omega(R, L, \varepsilon)$ is a strictly increasing function w.r.t. R and L .

Proof. To link the capacity to some quantity that can be optimized for neural networks, we need to apply the following theorem:

Theorem S3 ([S4], Theorem 11, Adapted). *With the neural network function space $\hat{\mathcal{F}}_{\text{NN}}$, let W be the total number of adjustable parameters, D the depth of the architecture. Let $\hat{\mathcal{G}} \subseteq \hat{\mathcal{F}}_{\text{NN}}$ be all functions into $[-R, R]^d$ representable on the architecture, and all these functions are at most L -Lipschitz. Then for all $0 < \varepsilon < 2R$,*

$$\mathcal{C}(\varepsilon, \hat{\mathcal{G}}, L^1) \leq \left(\frac{2e \cdot 2R \cdot DL}{\varepsilon} \right)^{2W}$$

Here, we need to prove firstly that the $\hat{\mathcal{F}}$ -dependent capacity of $\hat{\mathcal{G}}$ is bounded by a scaled independent capacity on L^1 of itself. We suppose that the MSE loss function (used in the definitions in Table S1) is bounded by some constant c . This is a reasonable assumption given that the input and output of neural networks are bounded in a compact set. Let us take G an $\frac{\varepsilon}{2\sqrt{c}}$ -cover of $\hat{\mathcal{G}}$ with L^1 -distance: $d_{L^1(\mathcal{P})}$ (see definition in Table S1). Therefore, for each $g \in \hat{\mathcal{G}}$ take $g' \in G$ such that $d_{L^1}(g, g') \leq \frac{\varepsilon}{2\sqrt{c}}$, then

$$\begin{aligned} d_{[\mathcal{P}, \hat{\mathcal{F}}]}(g, g') &= \int_{\mathcal{A} \times \mathcal{A}'} \sup_{f \in \hat{\mathcal{F}}} \left(\|(f+g)(x) - y\|_2^2 - \|(f+g')(x) - y\|_2^2 \right) d\mathcal{P}(x, y) \\ &\leq \int_{\mathcal{A} \times T\mathcal{A}} \sup_{f \in \hat{\mathcal{F}}} \|(g-g')(x)\|_2 (\|(f+g)(x) - y\|_2 + \|(f+g')(x) - y\|_2) d\mathcal{P}(x, y) \\ &\leq 2\sqrt{c} \int_{\mathbb{R}^d} \|(g-g')(x)\|_1 d\mathcal{P}(x) \leq \varepsilon \end{aligned}$$

Then we have the first inequality $\mathcal{C}_{\hat{\mathcal{F}}}(\varepsilon, \hat{\mathcal{G}}) \leq \mathcal{C}\left(\frac{\varepsilon}{2\sqrt{c}}, \hat{\mathcal{G}}, L^1\right)$. As we suppose that $\|g\|_\infty \leq R$ for all $g \in \hat{\mathcal{G}}$, then for all $g \in \hat{\mathcal{G}}$, we have $g(x) \in [-R, R]^d$. We now apply the Theorem S3 on $\hat{\mathcal{G}}$, we then have the following inequality

$$\log \mathcal{C}\left(\frac{\varepsilon}{2\sqrt{c}}, \hat{\mathcal{G}}, L^1\right) \leq 2W \log \frac{8e\sqrt{c}DRL}{\varepsilon} \quad (\text{S4})$$

where e is the base of the natural logarithm, W is the number of parameters of the architecture, D is the depth of the architecture. Then if we consider W, c, D as constants, the bound becomes:

$$\log \mathcal{C}\left(\frac{\varepsilon}{2\sqrt{c}}, \hat{\mathcal{G}}, L^1\right) \leq c_1 \log \frac{RL}{\varepsilon} + c_2 = \omega(R, L, \varepsilon) \quad (\text{S5})$$

for $c_1 = 2W$ and $c_2 = 2W \log 8e\sqrt{c}D$. \square

This leads us to choose for Ω a strictly increasing function that bounds $\omega(R, L, \varepsilon)$. Given the inequality (Eq. S3), this choice for Ω will allow us to bound practically the capacity of $\hat{\mathcal{G}}$.

Minimizing Ω will then reduce the effective capacity of the parametric set used to learn g_e . Concretely, we choose for Ω :

$$\Omega(g_e) = \|g_e\|_\infty^2 + \alpha \|g_e\|_{\text{Lip}}^2 \quad (7)$$

where $\alpha > 0$ is a hyper-parameter. This function is strictly convex and attains its unique minimum at the null function.

With this choice, let us instantiate Prop. 2 for our family of NNs. Let $r = \sup_{g \in \hat{\mathcal{G}}} \Omega(g)$, and $\omega(r, \varepsilon) = c_1 \log \frac{r}{\varepsilon\sqrt{\alpha}} + c_2$ (strictly increasing w.r.t. the r) for given parameters $c_1, c_2 > 0$. We have:

Table S2: Details for the results of evaluation error in test on linear systems in Fig. 1 .

Samples/env.	Method	$m = 1$	$m = 2$	$m = 4$	$m = 8$
$n = 2 \cdot K$	LEADS <i>no min.</i>	8.13±5.56 e-2	6.81±4.44 e-2	4.92±4.26 e-2	4.50±3.10 e-2
	LEADS (Ours)		5.11±3.20 e-2	3.93±2.88 e-2	2.10±0.96 e-2
$n = 4 \cdot K$	LEADS <i>no min.</i>	4.08±2.57 e-2	3.96±2.56 e-2	3.10±2.08 e-2	2.23±1.44 e-2
	LEADS (Ours)		2.74±1.96 e-2	1.61±1.24 e-2	1.02±0.74 e-2

Proposition S3. *If $r = \sup_{g \in \hat{\mathcal{G}}} \Omega(g)$ is finite, the number of samples n in Eq. 5, required to satisfy the error bound in Proposition 2 with the same $\delta, \varepsilon, \varepsilon_1$ and ε_2 becomes:*

$$n \geq \max \left\{ \frac{64}{\varepsilon^2} \left(\frac{1}{m} \log \frac{4\mathcal{C}_{\hat{\mathcal{G}}}(\frac{\varepsilon_1}{16}, \hat{\mathcal{F}})}{\delta} + \omega \left(r, \frac{\varepsilon_2}{16} \right) \right), \frac{16}{\varepsilon^2} \right\} \quad (\text{S6})$$

Proof. If $\Omega(g_e) \leq r$, we have $2 \log R \leq \log r$ and $2 \log L + \log \alpha \leq \log r$, then

$$\log RL \leq \log \frac{r}{\sqrt{\alpha}}$$

We can therefore bound $\omega(R, L, \varepsilon)$ by

$$\omega(R, L, \varepsilon) = c_1 \log \frac{RL}{\varepsilon} + c_2 \leq c_1 \log \frac{r}{\varepsilon \sqrt{\alpha}} + c_2 = \omega(r, \varepsilon)$$

The result follows from Proposition S2. \square

This means that the number of required samples will decrease with the size the largest possible $\Omega(g) = r$. The optimization process will reduce $\Omega(g_e)$ until a minimum is reached. The maximum size of the effective hypothesis space is then bounded and decreases throughout training. In particular, the following result follows:

Corollary S1. *Optimizing Eq. 4 for a given λ , we have that the number of samples n in Eq. 5 required to satisfy the error bound in Proposition 2 with the same $\delta, \varepsilon, \varepsilon_1$ and ε_2 is:*

$$n \geq \max \left\{ \frac{64}{\varepsilon^2} \left(\frac{1}{m} \log \frac{4\mathcal{C}_{\hat{\mathcal{G}}}(\frac{\varepsilon_1}{16}, \hat{\mathcal{F}})}{\delta} + \omega \left(\lambda \kappa, \frac{\varepsilon_2}{16} \right) \right), \frac{16}{\varepsilon^2} \right\} \quad (\text{S7})$$

where $\kappa = \sum_{e \in E} \sum_{i=1}^l \int_0^T \left\| \frac{dx_s^{e,i}}{dt} \right\|^2 ds$.

Proof. Denote $\mathcal{L}_\lambda(f, \{g_e\})$ the loss function defining Eq. 4. Consider a minimizer $(f^*, \{g_e^*\})$ of \mathcal{L}_λ . Then:

$$\mathcal{L}_\lambda(f^*, \{g_e^*\}) \leq \mathcal{L}_\lambda(0, \{0\}) = \kappa$$

which gives:

$$\forall e, \Omega(g_e^*) \leq \sum_e \Omega(g_e^*) \leq \lambda \kappa$$

Defining $\hat{\mathcal{G}} = \{g \in \hat{\mathcal{F}} \mid \Omega(g) \leq \lambda \kappa\}$, we then have that Eq. 4 is equivalent to:

$$\min_{f \in \hat{\mathcal{F}}, \{g_e\}_{e \in E} \in \hat{\mathcal{G}}} \sum_{e \in E} \left(\frac{\Omega(g_e)}{\lambda} + \sum_{i=1}^l \int_0^T \left\| \frac{dx_s^{e,i}}{dt} - (f + g_e)(x_s^{e,i}) \right\|^2 ds \right) \quad (\text{S8})$$

and the result follows from Proposition S3. \square

We can then decrease the sample complexity in the chosen NN family by: (a) increasing the number of training environments engaged in the framework, and (b) decreasing $\Omega(g_e)$ for all g_e , with $\Omega(g_e)$ instantiated as in Sec. 3.1. Ω provides a bound based on the largest output norm and the Lipschitz constant for a family of NNs. The experiments (Sec. 4) confirm that this is indeed an effective way to control the capacity of the approximating function family. Note that in our experiments, the number of samples needed in practice is much smaller than suggested by the theoretical bound.

Table S3: Detailed results of evaluation error in test on LV systems for Fig. 4. For the case of $m = 1$, all baselines except FT-RNN are equivalent to *One-Per-Env.*. The arrows indicate that the table cells share the same value.

Samples/env.	Method	$m = 1$	$m = 2$	$m = 4$	$m = 8$
$n = 1 \cdot K$	<i>One-For-All</i>	7.87±7.54 e-3	0.22±0.06	0.33±0.06	0.47±0.04
	<i>One-Per-Env.</i>	7.87±7.54 e-3	→		
	FT-RNN	4.02±3.17 e-2	1.62±1.14 e-2	1.62±1.40 e-2	1.08±1.03 e-2
	FT-NODE	7.87±7.54 e-3	7.63±5.84 e-3	4.18±3.77 e-3	4.92±4.19 e-3
	GBML-like	7.87±7.54 e-3	6.32±5.72 e-2	1.44±0.66 e-1	9.85±8.84 e-3
	<i>LEADS (Ours)</i>	7.87±7.54 e-3	3.65±2.99 e-3	2.39±1.83 e-3	1.37±1.14 e-3
$n = 2 \cdot K$	<i>One-For-All</i>	1.38±1.61 e-3	0.22±0.04	0.36±0.07	0.60±0.11
	<i>One-Per-Env.</i>	1.38±1.61 e-3	→		
	FT-RNN	7.20±7.12 e-2	2.72±4.00 e-2	1.69±1.57 e-2	1.38±1.25 e-2
	FT-NODE	1.38±1.61 e-3	9.02±8.81 e-3	1.11±1.05 e-3	1.00±0.95 e-3
	GBML-like	1.38±1.61 e-3	9.26±8.27 e-3	1.17±1.09 e-2	1.96±1.95 e-2
	<i>LEADS (Ours)</i>	1.38±1.61 e-3	8.65±9.61 e-4	8.40±9.76 e-4	6.02±6.12 e-4
$n = 4 \cdot K$	<i>One-For-All</i>	1.36±1.25 e-4	0.19±0.02	0.31±0.04	0.50±0.04
	<i>One-Per-Env.</i>	1.36±1.25 e-4	→		
	FT-RNN	8.69±8.36 e-4	3.39±3.38 e-4	3.02±1.50 e-4	2.26±1.45 e-4
	FT-NODE	1.36±1.25 e-4	1.74±1.65 e-4	1.78±1.71 e-4	1.39±1.20 e-4
	GBML-like	1.36±1.25 e-4	2.57±7.18 e-3	2.65±3.26 e-3	2.36±3.58 e-3
	<i>LEADS (Ours)</i>	1.36±1.25 e-4	1.10±0.92 e-4	1.03±0.98 e-4	9.66±9.79 e-5
$n = 8 \cdot K$	<i>One-For-All</i>	5.98±5.13 e-5	0.16±0.03	0.35±0.06	0.52±0.06
	<i>One-Per-Env.</i>	5.98±5.13 e-5	→		
	FT-RNN	2.09±1.73 e-4	1.18±1.16 e-4	1.13±1.13 e-4	9.13±8.31 e-5
	FT-NODE	5.98±5.13 e-5	6.91±4.46 e-5	7.82±6.95 e-5	6.88±6.39 e-5
	GBML-like	5.98±5.13 e-5	1.02±1.68 e-4	1.41±2.68 e-4	0.99±1.53 e-4
	<i>LEADS (Ours)</i>	5.98±5.13 e-5	5.47±4.63 e-5	4.52±3.98 e-5	3.94±3.49 e-5

Table S4: Test MSE of experiments on LV ($m = 4, n = 1 \cdot K$) with different empirical norms.

Empirical Norm	$p = 1$	$p = 2$	$p = 3$	$p = 10$	$p = \infty$
Test MSE	2.30e-3	2.36e-3	2.34e-3	3.41e-3	6.12e-3

C Optimizing Ω in practice

In Sec. 3.3, we developed an instantiation of the *LEADS* framework for neural networks. We proposed to control the capacity of the g_e s components through a penalization function Ω defined as $\Omega(g_e) = \|g_e\|_\infty^2 + \alpha \|g_e\|_{\text{Lip}}^2$. This definition ensures the properties required to control the sample complexity.

However, in practice, both terms in $\Omega(g_e)$ are difficult to compute as they do not yield an analytical form for neural networks. For a fixed activation function, the Lipschitz-norm of a trained model only depends on the model parameters and, for our class of neural networks, can be bounded by the spectral norms of the weight matrices, as described in Sec. 4.4. This allows for a practical implementation.

The infinity norm on its side depends on the domain definition of the function and practical implementations require an empirical estimate. Since there is no trivial estimator for the infinity norm of a function, we performed tests with different proxies such as the *empirical* L^p and L^∞ norms, respectively defined as $\|g_e\|_{L^p(\hat{\mathcal{F}}_e)} = \left(\frac{1}{n} \sum_{x \in \hat{\mathcal{F}}_e} |g_e(x)|^p \right)^{1/p}$ for $1 \leq p < \infty$ and $\|g_e\|_{L^\infty(\hat{\mathcal{F}}_e)} = \max_{x \in \hat{\mathcal{F}}_e} |g_e(x)|$. Here $|\cdot|$ is an ℓ^2 vector norm. Note that on a finite set of points, these norms reduce to vector norms $\|(|g_e(x_1)|, \dots, |g_e(x_n)|)\|_p$. They are then all equivalent on the space defined by the training set. Table S4 shows the results of experiments performed on LV equation with different $1 \leq p \leq \infty$. Overall we found that L^p for small values of p worked better and chose in our experiments set $p = 2$.

Moreover, using both minimized quantities $\|g_e\|_{L^2(\hat{\mathcal{P}}_e)}^2$ and the spectral norm of the product of weight matrices, denoted $L(g_e)$ and $\Pi(g_e)$ respectively, we can give a bound on $\Omega(g_e)$. First, for any x in the compact support of \mathcal{P}_e , we have that, fixing some $x_0 \in \hat{\mathcal{P}}_e$:

$$|g_e(x)| \leq |g_e(x) - g_e(x_0)| + |g_e(x_0)|$$

For the first term:

$$|g_e(x) - g_e(x_0)| \leq \|g_e\|_{\text{Lip}}|x - x_0| \leq \Pi(g_e)|x - x_0|$$

and the support of \mathcal{P}_e being compact by hypothesis, denoting by δ its diameter:

$$|g_e(x) - g_e(x_0)| \leq \delta \Pi(g_e)$$

Moreover, for the second term:

$$|g_e(x_0)| = \sqrt{|g_e(x_0)|^2} \leq \sqrt{L(g_e)}$$

and summing both contributions gives us the bound:

$$\|g_e\|_\infty \leq \delta \Pi(g_e) + \sqrt{L(g_e)}$$

so that:

$$\Omega(g_e) \leq (\delta + \alpha)\Pi(g_e) + \sqrt{L(g_e)}$$

Note that this estimation is a crude one and improvements can be made by considering the closest x_0 from x and taking δ to be the maximal distance between points not from the support of \mathcal{P}_e and $\hat{\mathcal{P}}_e$.

Finally, we noticed that minimizing $\|\frac{g_e}{id}\|_{L^2(\hat{\mathcal{P}}_e)}^2$ in domains bounded away from zero gave better results as normalizing by the norm of the output allowed to adaptively rescale the computed norm. Formally, minimizing this quantity does not fundamentally change the optimization as we have that:

$$\forall x, \frac{1}{M^2}|g_e(x)|^2 \leq \left| \frac{g_e(x)}{x} \right|^2 \leq \frac{1}{m^2}|g_e(x)|^2$$

meaning that:

$$\frac{1}{M^2}L(g_e) \leq \left\| \frac{g_e}{id} \right\|_{L^2(\hat{\mathcal{P}}_e)}^2 \leq \frac{1}{m^2}L(g_e)$$

where m, M are the lower and upper bound of $|x|$ on the support of \mathcal{P}_e with $m > 0$ by hypothesis (the quantity we minimize is still higher than $L(g_e)$ even if this is not the case).

D Additional experimental details

D.1 Details on the environment dynamics

Lotka-Volterra (LV). The model dynamics follow the ODE:

$$\frac{du}{dt} = \alpha u - \beta uv, \frac{dv}{dt} = \delta uv - \gamma v$$

with u, v the number of prey and predator, $\alpha, \beta, \gamma, \delta > 0$ defining how the two species interact. The initial conditions u_0^i, v_0^i are sampled from a uniform distribution $P_0 = \text{Unif}([1, 2]^2)$. We characterize the dynamics by $\theta = (\alpha/\beta, \gamma/\delta) \in \Theta = \{0.5, 1, 1.44, 1.5, 1.86, 2\}^2$. An environment e is then defined by parameters θ_e sampled from a uniform distribution over the parameter set Θ .

Gray-Scott (GS). The governing PDE is:

$$\frac{\partial u}{\partial t} = D_u \Delta u - uv^2 + F(1 - u), \frac{\partial v}{\partial t} = D_v \Delta v + uv^2 - (F + k)v$$

where the u, v represent the concentrations of two chemical components in the spatial domain S with periodic boundary conditions. D_u, D_v denote the diffusion coefficients respectively for u, v , and are held constant to $D_u = 0.2097, D_v = 0.105$, and F, k are the reaction parameters depending on the environment. As for the initial conditions $(u_0, v_0) \sim P_0$, we place 3 2-by-2 squares at uniformly sampled positions in S to trigger the reactions. The values of (u_0, v_0) are fixed to $(0, 1)$ outside the squares and to $(1 - \epsilon, \epsilon)$ with a small $\epsilon > 0$ inside. An environment e is defined by its parameters $\theta_e = (F_e, k_e) \in \Theta = \{(0.037, 0.060), (0.030, 0.062), (0.039, 0.058)\}$. We consider a set of θ_e parameters uniformly sampled from the environment distribution Q on Θ .

Navier-Stokes (NS). We consider the Navier-Stokes PDE for incompressible flows:

$$\frac{\partial w}{\partial t} = -v \cdot \nabla w + \nu \Delta w + \xi \quad \nabla \cdot v = 0$$

where v is the velocity field, $w = \nabla \times v$ is the vorticity, both v, w lie in a spatial domain S with periodic boundary conditions, ν is the viscosity and ξ is the constant forcing term in the domain S . We fix $\nu = 10^{-3}$ across the environments. We sample the initial conditions $w_0^e \sim P_0$ as in [S6]. An environment e is defined by its forcing term $\xi_e \in \Theta_\xi = \{\xi_1, \xi_2, \xi_3, \xi_4\}$ with

$$\begin{aligned} \xi_1(x, y) &= 0.1(\sin(2\pi(x + y)) + \cos(2\pi(x + y))) \\ \xi_2(x, y) &= 0.1(\sin(2\pi(x + y)) + \cos(2\pi(x + 2y))) \\ \xi_3(x, y) &= 0.1(\sin(2\pi(x + y)) + \cos(2\pi(2x + y))) \\ \xi_4(x, y) &= 0.1(\sin(2\pi(2x + y)) + \cos(2\pi(2x + y))) \end{aligned}$$

where $(x, y) \in S$ is the position in the domain S . We uniformly sampled a set of forcing terms from Q on Θ_ξ .

Linear ODE. We take an example of linear ODE expressed by the following formula:

$$\frac{du_t}{dt} = L_{Q\Lambda Q^\top}(u_t) = Q\Lambda Q^\top u_t$$

where $u_t \in \mathbb{R}^8$ is the system state, $Q \in M_{8,8}(\mathbb{R})$ is an orthogonal matrix such that $QQ^\top = 1$, and $\Lambda \in M_{8,8}(\mathbb{R})$ is a diagonal matrix containing eigenvalues. We sample Λ_e from a uniform distribution on $\Theta_\Lambda = \{\Lambda_1, \dots, \Lambda_8\}$, defined for each Λ_i by:

$$[\Lambda_i]_{jj} = \begin{cases} 0, & \text{if } i = j \text{ for } i, j \in \{1, \dots, 8\}, \\ -0.5, & \text{otherwise.} \end{cases}$$

which means that the i -th eigenvalue is set to 0, while others are set to a common value -0.5 .

D.2 Choosing hyperparameters

As usual, the hyperparameters need to be tuned for each considered set of systems. We therefore chose the hyperparameters using standard cross-validation techniques. We did not conduct a systematic sensitivity analysis. In practice, we found that: (a) if the regularization term is too large w.r.t. the trajectory loss, the model cannot fit the trajectories, and (b) if the regularization term is too small, the performance is similar to *LEADS no min*. The candidate hyperparameters are defined on a very sparse grid, for example, for neural nets, $(10^3, 10^4, 10^5, 10^6)$ for λ and $(10^{-2}, 10^{-3}, 10^{-4}, 10^{-5})$ for α .

D.3 Details on the experiments with a varying number of environments

We conducted large-scale experiments respectively for linear ODEs (Sec. 3.2, Fig. 1) and LV (Sec. 4, Fig. 4) to compare the tendency of *LEADS* w.r.t. the theoretical bound and the baselines by varying the number of environments available for the instantiated model.

To guarantee the comparability of the test-time results, we need to use the same test set when varying the number of environments. We therefore propose to firstly generate a global set of environments, separate it into subgroups for training, then we test these separately trained models on the global test set.

We performed the experiments as follows:

- In the training phase, we consider $M = 8$ environments in total in the environment set E_{total} . We denote here the cardinality of an environment set E by $\text{card}(E)$, the environments are then arranged into $b = 1, 2, 4$ or 8 disjoint groups of the same size, i.e. $\{E_1, \dots, E_b\}$ such that $\bigcup_{i=1}^b E_i = E_{\text{total}}$, $\text{card}(E_1) = \dots = \text{card}(E_b) = \lfloor M/b \rfloor =: m$, where m is the number of environments per group, and $E_i \cap E_j = \emptyset$ whenever $i \neq j$. For example, for $m = 1$, all the original environments are gathered into one global environment, when for $m = 8$ we keep all the original environments. The methods are then instantiated respectively for each E_i . For example, for *LEADS* with b environment groups, we instantiate $LEADS_1, \dots, LEADS_b$ respectively on E_1, \dots, E_b . Other frameworks are applied in the same way.

Table S5: Results on 2 novel environments for LV, GS, and NS at different training steps with n data points per env. The arrows indicate that the table cells share the same value.

Dataset	Training Schema	Test MSE at training step		
		50	2500	10000
LV ($n = 1 \cdot K$)	<i>Pre-trained-f-Only</i>	0.36	—————→	
	<i>One-Per-Env.</i> from scratch	0.23	8.85e-3	3.05e-3
	<i>Pre-trained-f-Plus-Trained-g_e</i>	0.73	1.36e-3	1.11e-3
GS ($n = 1 \cdot K$)	<i>Pre-trained-f-Only</i>	5.44e-3	—————→	
	<i>One-Per-Env.</i> from scratch	4.20e-2	5.53e-3	3.05e-3
	<i>Pre-trained-f-Plus-Trained-g_e</i>	2.29e-3	1.45e-3	1.27e-3
NS ($n = 8 \cdot K$)	<i>Pre-trained-f-Only</i>	1.75e-1	—————→	
	<i>One-Per-Env.</i> from scratch	6.76e-2	1.70e-2	1.18e-2
	<i>Pre-trained-f-Plus-Trained-g_e</i>	1.37e-2	8.07e-3	7.14e-3

Note that when $m = 1$, having $b = 8$ environment groups of one single environment, *One-For-All*, *One-Per-Env.* and *LEADS* are reduced to *One-Per-Env.* applied on all M environments. We can see in Fig. 4 that each group of plots starts from the same point.


- In the test phase, the performance of the model trained with the group E_i is tested with the test samples of the corresponding group. Then we take the mean error over all b groups to obtain the results on all M environments. Note that the result at each point in figures 1 and 4 is calculated on the same total test set, which guarantees the comparability between results.

D.4 Additional experimental results


Experiments with a varying number of environments We show in tables S2 and S3 the detailed results used for the plots in figures 1 and 4, compared to baseline methods.

Learning in novel environments We conducted same experiments as in Sec. 4.3 to learn in unseen environments for GS and NS datasets. The test MSE at different training steps is shown in Table S5.


Full-length trajectories We provide in figures S1-S4 the full-length sample trajectories for GS and NS of Fig. 2.




(a) *One-Per-Env.*



(b) FT-NODE




(c) *LEADS*




(d) Ground truth


Figure S1: Full-length prediction comparison of Fig. 2 for GS. In each figure, from top to bottom, the trajectory snapshots are output respectively from 3 training environments. The temporal resolution of each sequence is $\Delta t = 40$.



(a) Difference between *One-Per-Env.* and Ground truth




(b) Difference between FT-NODE and Ground truth




(c) Difference between *LEADS* and Ground truth


Figure S2: Full-length error maps of Fig. 2 for GS. In each figure, from top to bottom, the trajectory snapshots correspond to 3 training environments, one per row. The temporal resolution of each sequence is $\Delta t = 40$.




(a) *One-Per-Env.*



(b) FT-NODE




(c) *LEADS*




(d) Ground truth


Figure S3: Full-length prediction comparison of Fig. 2 for NS. In each figure, from top to bottom, the trajectory snapshots correspond to 3 training environments. The temporal resolution of each sequence is $\Delta t = 1$.



(a) Difference between *One-Per-Env.* and Ground truth



(b) Difference between FT-NODE and Ground truth



(c) Difference between *LEADS* and Ground truth

Figure S4: Full-length error maps of Fig. 2 for NS. In each figure, from top to bottom, the trajectory snapshots correspond to from 3 training environments. The temporal resolution of each sequence is $\Delta t = 1$.

References

- [S1] P. L. Bartlett, D. J. Foster, and M. J. Telgarsky. Spectrally-normalized margin bounds for neural networks. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30, pages 6240–6249. Curran Associates, Inc., 2017.
- [S2] J. Baxter. A model of inductive bias learning. *J. Artif. Int. Res.*, 12(1):149–198, Mar. 2000.
- [S3] L. Chizat and F. Bach. On the global convergence of gradient descent for over-parameterized models using optimal transport. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31, pages 3036–3046. Curran Associates, Inc., 2018.
- [S4] D. Haussler. Decision theoretic generalizations of the pac model for neural net and other learning applications. *Information and Computation*, 100(1):78 – 150, 1992.
- [S5] P. Kidger and T. Lyons. Universal Approximation with Deep Narrow Networks. In J. Abernethy and S. Agarwal, editors, *Proceedings of Thirty Third Conference on Learning Theory*, volume 125 of *Proceedings of Machine Learning Research*, pages 2306–2327. PMLR, 09–12 Jul 2020.
- [S6] Z. Li, N. B. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart, and A. Anandkumar. Fourier neural operator for parametric partial differential equations. In *International Conference on Learning Representations*, 2021.
- [S7] S. Shalev-Shwartz and S. Ben-David. *Covering Numbers*, page 337–340. Cambridge University Press, 2014.