LEARNING WHERE TO LEARN: TRAINING DATA DISTRIBUTION OPTIMIZATION FOR SCIENTIFIC MACHINE LEARNING

Anonymous authors

Paper under double-blind review

ABSTRACT

In scientific machine learning, models are routinely deployed with parameter values or boundary conditions far from those used in training. This paper studies the learning-where-to-learn problem of designing a training data distribution that minimizes average prediction error across a family of deployment regimes. A theoretical analysis shows how the training distribution shapes deployment accuracy. This motivates two adaptive algorithms based on bilevel or alternating optimization in the space of probability measures. Discretized implementations using parametric distribution classes or nonparametric particle-based gradient flows deliver optimized training distributions that outperform nonadaptive designs. Once trained, the resulting models exhibit improved sample complexity and robustness to distribution shift. This framework unlocks the potential of principled data acquisition for learning functions and solution operators of partial differential equations.

1 Introduction

In scientific settings, machine learning models encounter data at deployment time that can differ substantially from the data seen during training. Standard empirical risk minimization (ERM), which minimizes an average loss over the fixed training distribution, may yield models that perform well in-sample yet degrade under distribution shift. This challenge is acute in operator learning for partial differential equations (PDEs) and infinite-dimensional inverse problems. For example, in data-driven electrical impedance tomography (EIT), neural operators are trained to map electrical patterns on the boundary to internal conductivities (Fan & Ying, 2020; Molinaro et al., 2023; Park et al., 2021). The input is a collection of voltage and corresponding current pairs that represent Dirichlet and Neumann boundary conditions, respectively. The probability distribution ν that boundary voltages are sampled from is a user choice. Crucially, the choice of ν shapes how well the trained model generalizes to new boundary measurement distributions. Fig. 1 illustrates this effect with a variant of the Neural Inverse Operator (NIO) architecture (Molinaro et al., 2023); see **SM** E–F for details.

These observations motivate the central premise of the present work: learning where to learn by designing an optimal training distribution ν for a prescribed family of deployment regimes. Guided by theoretical and practical considerations, this paper develops a principled and computationally feasible framework for optimizing ν in function regression and operator learning tasks. It establishes intelligent data acquisition as a crucial component of scientific machine learning (SciML) workflows.

Contributions. Our main contributions are as follows.

- (C1) Lipschitz-based distribution shift bounds. We develop quantitative bounds that connect out-of-distribution (OOD) error to the training error, Lipschitz continuity of the model and target map, and the Wasserstein mismatch between training and test distributions.
- (C2) Algorithms for training distribution optimization. Motivated by the theory, we develop two adaptive algorithms that actively optimize the training distribution for general model classes: (1) a bilevel procedure that directly minimizes OOD error and (2) a computationally efficient alternating scheme that minimizes an upper bound on the OOD error. We implement (1) and (2) over the space of probability measures by considering (i) parametric training distribution families (e.g., Gaussian processes) and (ii) nonparametric particle representations of the training distribution (e.g., that are updated with Wasserstein gradient flows).

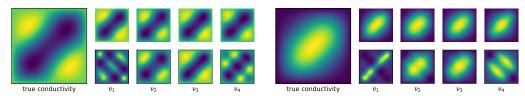


Figure 1: Two conductivity samples in EIT. The true conductivity is on the left of each panel, followed by predictions from NIO models trained on four Dirichlet boundary condition distributions $\{\nu_i\}_{i=1}^4$. Top row: in-distribution predictions; bottom row: out-of-distribution predictions. See SM E-F for details.

(C3) Empirical performance in SciML tasks. We perform numerical experiments involving function approximation as well as forward and inverse operator learning for a range of elliptic and hyperbolic PDEs (EIT, Darcy flow, and radiative transport). Consistent with our theory, the results show that the training data distribution strongly governs OOD deployment error. In the experiments, the two proposed adaptive algorithms efficiently optimize the training distribution, greatly reduce OOD error, and often outperform nonadaptive baselines.

Related work. The careful selection of sampling points for regression appears frequently in numerical analysis (Barron et al., 2008; Krieg et al., 2022). Weighted least-squares regression with Christoffel sampling is one prominent example (Cohen & Migliorati, 2017; Adcock et al., 2022). One of the main findings of these works is that training on samples from the target test distribution is a suboptimal strategy in general. Instead, a change of measure is derived by minimizing an upper bound on average error. These ideas have their origin in numerical integration and Monte Carlo methods (Caflisch, 1998). Some other related design strategies include greedy Bayesian herding (Huszár & Duvenaud, 2012) and kernel quadrature (Epperly & Moreno, 2023). However, quadrature-based sampling does not necessarily deliver optimal performance for regression problems.

In SciML for PDEs, several adaptive sampling and reweighting strategies have emerged, particularly for unsupervised methods such as physics-informed neural networks (PINNs) (Wu et al., 2023). Common approaches either reweight a fixed pool of training data samples (Chen et al., 2025a; McClenny & Braga-Neto, 2023) or sample new data points from a density proportional to the PDE residual (Gao et al., 2023; Tang et al., 2023). Similar methods based on importance sampling, reweighting, or coreset selection have been applied to neural operators (Wang et al., 2022b; Toscano et al., 2025). Others use residual information to move samples directly (Ouyang et al., 2025) or iteratively expand training datasets (Lye et al., 2021). The current work builds on these advances by operating at the level of probability distributions. This enables a wide range of parametrizations and discretizations beyond samples only and provides extra freedom to explore the design space.

For a discussion of broader connections to other areas of machine learning and statistics, see SM A.

Outline. This paper is organized as follows. Sec. 2 provides the necessary background. Sec. 3 establishes motivating theoretical results (C1). Sec. 4 proposes two implementable data distribution selection algorithms (C2). Sec. 5 performs numerical studies (C3) and Sec. 6 gives concluding remarks. Proof and implementation details are deferred to the supplementary material (SM).

2 Preliminaries

This section sets the notation, provides essential background on key mathematical concepts, and reviews the framework of operator learning. **SM** B contains additional reference material.

Notation. This paper works with an abstract input space $(\mathcal{U}, \langle \cdot, \cdot \rangle_{\mathcal{U}}, \| \cdot \|_{\mathcal{U}})$ and output space $(\mathcal{Y}, \langle \cdot, \cdot \rangle_{\mathcal{Y}}, \| \cdot \|_{\mathcal{Y}})$, both assumed to be real separable Hilbert spaces. We often drop the subscripts on norms and inner products when the meaning is clear from the context. Unless otherwise stated, \mathcal{U} or \mathcal{Y} could be infinite-dimensional. We denote the Lipschitz constant of a map $F: \mathcal{U} \to \mathcal{Y}$ by $\operatorname{Lip}(F) := \sup_{u \neq u'} \|F(u) - F(u')\|_{\mathcal{Y}} / \|u - u'\|_{\mathcal{U}}$. For any element ζ of a metric space, the Dirac probability measure assigning unit mass to ζ is written as δ_{ζ} . For $p \in [1, \infty]$, (Bochner) $L^p_{\mu}(\mathcal{U}; \mathcal{Y})$ spaces are defined in the usual way. We say that a probability measure $\mu \in \mathscr{P}(\mathcal{U})$ belongs to $\mathscr{P}_p(\mathcal{U})$ if $\mathsf{m}_p(\mu) < \infty$, where $\mathsf{m}_p: \mathscr{P}(\mathcal{U}) \to \mathbb{R}_{\geq 0} \cup \{\infty\}$ denotes the uncentered p-th moment defined by $\mu \mapsto \mathsf{m}_p(\mu) := \mathbb{E}_{u \sim \mu} \|u\|_{\mathcal{U}}^p$. We often work with Gaussian measures $\mathcal{N}(m, \mathcal{C}) \in \mathscr{P}_2(\mathcal{U})$, where

 $m \in \mathcal{U}$ is the mean and $\mathcal{C}: \mathcal{U} \to \mathcal{U}$ is the symmetric trace-class positive semidefinite covariance operator; see (Dashti & Stuart, 2017, App. A.3) for details in the infinite-dimensional setting.

Background. We review several key concepts that frequently appear in this work.

The *pushforward operator* maps a measure from one measurable space to another via a measurable map. Specifically, given measure spaces (X, \mathcal{B}_X, μ) and (Y, \mathcal{B}_Y, ν) and a measurable map $T \colon X \to Y$, the pushforward operator $T_\#$ assigns to μ a measure $T_\#\mu$ on Y defined by $T_\#\mu(B) = \mu(T^{-1}(B))$ for all $B \in \mathcal{B}_Y$, where $T^{-1}(B) = \{x \in X \mid T(x) \in B\}$ is the pre-image of B. Intuitively, $T_\#$ redistributes the measure μ on X to a measure $T_\#\mu$ on Y following the structure of the map T.

Definition 2.1 (p-Wasserstein distance). Consider probability measures μ and ν on a metric space (X,d) with finite p-th moments, where $p\in[1,\infty)$. Define the set of all couplings $\Gamma(\mu,\nu)\coloneqq\left\{\gamma\in\mathscr{P}(X\times X)\,\middle|\,\pi_\#^1\gamma=\mu,\,\pi_\#^2\gamma=\nu\right\}$, where $\pi^1(x,y)=x$ and $\pi^2(x,y)=y$. The p-Wasserstein distance between μ and ν is $\mathsf{W}_p(\mu,\nu)\coloneqq(\inf_{\gamma\in\Gamma(\mu,\nu)}\int_{X\times X}\mathsf{d}(x,y)^p\,\gamma(dx,dy))^{1/p}$.

The cases of p=1 and p=2 are often considered, corresponding to the W_1 and W_2 metrics (Villani, 2021). The case when μ and ν are both Gaussian and X is a Hilbert space is also used frequently.

We will use the notion of random measure (Kallenberg, 2017) to quantify OOD error in Secs. 3-4.

Definition 2.2 (random measure). Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space, and let (M, \mathcal{B}) be a measurable space. A *random measure* on M is a map $\Xi \colon \Omega \to \mathscr{P}(M)$, $\Xi(\omega) := \mu_{\omega}$, such that for every $A \in \mathcal{B}$, the function $\omega \mapsto \mu_{\omega}(A)$ is a random variable, i.e., measurable from (Ω, \mathcal{F}) to $(\mathbb{R}_{\geq 0}, \mathcal{B}(\mathbb{R}))$.

Operator learning. Operator learning provides a powerful framework for approximating mappings $\mathcal{G}^* \colon \mathcal{U} \to \mathcal{Y}$ between function spaces. It plays an increasingly pivotal role for accelerating forward and inverse problems in the computational sciences (Boullé & Townsend, 2024; Kovachki et al., 2023; 2024b; Nelsen & Stuart, 2024). As in traditional regression, the learned operator $\mathcal{G}_{\widehat{\theta}}$ is parametrized over a set Θ and minimizes an *empirical risk*

$$\widehat{\theta} \in \underset{\theta \in \Theta}{\operatorname{arg\,min}} \, \frac{1}{N} \sum_{n=1}^{N} \left\| \mathcal{G}^{\star}(u_n) - \mathcal{G}_{\theta}(u_n) \right\|_{\mathcal{Y}}^{2}. \tag{2.1}$$

The Hilbert norm in (2.1) and discretization-invariant architectures (e.g., DeepONet (Lu et al., 2021), Fourier Neural Operator (FNO) (Li et al., 2021), NIO (Molinaro et al., 2023)) are points of departure from classical function approximation. In (2.1), $\{u_n\} \sim \nu^{\otimes N}$, where ν is the training distribution over input functions. The choice and properties of ν are crucial. An inadequately selected ν may result in $\mathcal{G}_{\widehat{\theta}}(u)$ being a poor approximation to $\mathcal{G}^{\star}(u)$ for $u \sim \nu'$, where $\nu' \neq \nu$ (Adcock et al., 2022; Boullé et al., 2023; de Hoop et al., 2023; Li et al., 2024; Musekamp et al., 2024; Pickering et al., 2022; Subedi & Tewari, 2024). The present work aims to improve accuracy under distribution shift in both function approximation and operator learning settings.

3 LIPSCHITZ THEORY FOR OUT-OF-DISTRIBUTION ERROR BOUNDS

This section studies how loss functions behave under distribution shifts and the implications of such shifts on average-case OOD accuracy. The assumed Lipschitz continuity of the target map and the elements in the approximation class play a central role in the analysis.

Distribution shift inequalities. In the infinite data limit, the functional in (2.1) naturally converges to the so-called *expected risk*, which serves as a notion of model accuracy in this work. The following result forms the foundation for our forthcoming data selection algorithms in Sec. 4.

Proposition 3.1 (distribution shift error). Let \mathcal{G}_1 and \mathcal{G}_2 both be Lipschitz continuous maps from Hilbert space \mathcal{U} to Hilbert space \mathcal{Y} . For any $\nu \in \mathscr{P}_1(\mathcal{U})$ and $\nu' \in \mathscr{P}_1(\mathcal{U})$, it holds that

$$\mathbb{E}_{u' \sim \nu'} \| \mathcal{G}_1(u') - \mathcal{G}_2(u') \|_{\mathcal{Y}} \leq \mathbb{E}_{u \sim \nu} \| \mathcal{G}_1(u) - \mathcal{G}_2(u) \|_{\mathcal{Y}} + \left(\mathrm{Lip}(\mathcal{G}_1) + \mathrm{Lip}(\mathcal{G}_2) \right) W_1(\nu, \nu') . \quad (3.1)$$

Moreover, for any $\mu \in \mathscr{P}_2(\mathcal{U})$ and $\mu' \in \mathscr{P}_2(\mathcal{U})$, it holds that

$$\mathbb{E}_{u' \sim \mu'} \| \mathcal{G}_1(u') - \mathcal{G}_2(u') \|_{\mathcal{Y}}^2 \leq \mathbb{E}_{u \sim \mu} \| \mathcal{G}_1(u) - \mathcal{G}_2(u) \|_{\mathcal{Y}}^2 + c(\mathcal{G}_1, \mathcal{G}_2, \mu, \mu') \mathsf{W}_2(\mu, \mu') , \qquad (3.2)$$
where $c(\mathcal{G}_1, \mathcal{G}_2, \mu, \mu')$ equals

$$\left(\operatorname{Lip}(\mathcal{G}_1) + \operatorname{Lip}(\mathcal{G}_2)\right) \sqrt{4 \left(\operatorname{Lip}(\mathcal{G}_1) + \operatorname{Lip}(\mathcal{G}_2)\right)^2 \left[\mathsf{m}_2(\mu) + \mathsf{m}_2(\mu')\right] + 16 \left[\|\mathcal{G}_1(0)\|_{\mathcal{Y}}^2 + \|\mathcal{G}_2(0)\|_{\mathcal{Y}}^2\right]}.$$

See SM C for the proof. Provided that the maps \mathcal{G}_1 and \mathcal{G}_2 are Lipschitz, Prop. 3.1 shows that the functional $\mu \mapsto \mathbb{E}_{u \sim \mu} \|\mathcal{G}_1(u) - \mathcal{G}_2(u)\|$ is globally Lipschitz continuous in the W₁ metric. Similarly, the functional $\mu \mapsto \mathbb{E}_{u \sim \mu} \|\mathcal{G}_1(u) - \mathcal{G}_2(u)\|^2$ is *locally* Lipschitz continuous in the W₂ metric; in particular, it is Lipschitz on sets with uniformly bounded second moments. We remark that the Lipschitz map hypotheses in Prop. 3.1 are invoked for simplicity. Hölder continuity is also compatible with Wasserstein distances, and the preceding result can be adapted to hold under such regularity.

Average-case performance. OOD performance may be quantified in several ways. This paper chooses *average-case accuracy* with respect to the law $\mathbb{Q} \in \mathscr{P}(\mathscr{P}(\mathcal{U}))$ of a random probability measure $\nu' \sim \mathbb{Q}$. For a ground truth map \mathcal{G}^* and any hypothesis \mathcal{G} , we strive for

$$\mathscr{E}_{\mathbb{Q}}(\mathcal{G}) := \mathbb{E}_{\nu' \sim \mathbb{Q}} \, \mathbb{E}_{u' \sim \nu'} \| \mathcal{G}^{\star}(u') - \mathcal{G}(u') \|_{\mathcal{V}}^{2} \tag{3.3}$$

to be small. An equivalent viewpoint in (D.2) from **SM** D.1 rewrites (3.3) as a single expectation against a mixture. Other approaches to quantify OOD error are possible, such as worst-case analysis.

The following corollary is useful. It is a consequence of Prop. 3.1 and the Cauchy–Schwarz inequality.

Corollary 3.2 (basic inequality). For any $\mu \in \mathscr{P}_2(\mathcal{U})$ and maps \mathcal{G}^* and \mathcal{G} , it holds that

$$\mathscr{E}_{\mathbb{Q}}(\mathcal{G}) \leq \mathbb{E}_{u \sim \mu} \| \mathcal{G}^{\star}(u) - \mathcal{G}(u) \|_{\mathcal{Y}}^{2} + \sqrt{\mathbb{E}_{\nu' \sim \mathbb{Q}} c^{2}(\mathcal{G}^{\star}, \mathcal{G}, \mu, \nu')} \sqrt{\mathbb{E}_{\nu' \sim \mathbb{Q}} \mathsf{W}_{2}^{2}(\mu, \nu')} \,. \tag{3.4}$$

The upper bound (3.4) on the average-case OOD error $\mathscr{E}_{\mathbb{Q}}$ is intuitive. We view the arbitrary probability measure μ as the training data distribution that we are free to select. The first term on the right-hand side of (3.4) then represents the in-distribution generalization error of \mathcal{G} . On the other hand, the second term can be considered a regularizer that penalizes how far μ is from \mathbb{Q} on average with respect to the W_2 metric. The factor involving the quantity c further penalizes the size of moments of μ and \mathbb{Q} as well as the size of the Lipschitz constants of \mathcal{G}^* and \mathcal{G} .

We instantiate Cor. 3.2 in the concrete context of Gaussian mixtures in SM C, Ex. C.1.

Although the preceding OOD results hold in a quite general setting, they do not distinguish between an arbitrary model and a trained model obtained by ERM (2.1). Sharpening the bounds for a trained model would require probabilistic estimates on the model's Lipschitz constant and in-distribution error. These are challenging theoretical tasks. Instead, we turn our attention to implementable algorithms inspired by Cor. 3.2 that are designed to reduce OOD error in practice.

4 PRACTICAL ALGORITHMS FOR TRAINING DISTRIBUTION DESIGN

Let $\mathcal{G}^{\star} \colon \mathcal{U} \to \mathcal{Y}$ be the target mapping and \mathbb{Q} be a probability measure over $\mathscr{P}_2(\mathcal{U})$. We assume query access to an oracle that returns the label $\mathcal{G}^{\star}(u) \in \mathcal{Y}$ when presented with an input $u \in \mathcal{U}$. The idealized infinite data training distribution selection problem that we wish to address is

$$\inf_{\nu \in \mathscr{P}_{2}(\mathcal{U})} \left\{ \mathscr{E}_{\mathbb{Q}}(\widehat{\mathcal{G}}^{(\nu)}) \, \middle| \, \widehat{\mathcal{G}}^{(\nu)} \in \operatorname*{arg\,min}_{\mathcal{G} \in \mathcal{H}} \mathbb{E}_{u \sim \nu} \big\| \mathcal{G}^{\star}(u) - \mathcal{G}(u) \big\|_{\mathcal{Y}}^{2} \right\}, \tag{4.1}$$

where $\mathscr{E}_{\mathbb{Q}}$ is the average-case OOD accuracy functional (3.3) and \mathcal{H} is the model hypothesis class.

Eqn. (4.1) is a bilevel optimization problem. The inner minimization problem seeks the best-fit model $\widehat{\mathcal{G}}^{(\nu)}$ for a fixed training distribution ν . The outer problem minimizes $\mathscr{E}_{\mathbb{Q}}(\widehat{\mathcal{G}}^{(\nu)})$ over ν in the space of probability measures. The map $\nu \mapsto \widehat{\mathcal{G}}^{(\nu)}$ can be highly nonlinear, which poses a challenge to solve (4.1). To devise practical solution methods, this section considers the bilevel optimization problem (4.1) in Sec. 4.1 and a computationally tractable relaxation of (4.1) in Sec. 4.2.

Remark 4.1 (finite data). It is necessary to use finite data to estimate the optimal training distribution (SM D.1). We thus assume sample access to the chosen feasible set of training distributions, e.g., i.i.d. sampling or interacting particles. Moreover, we only require sample access to samples from \mathbb{Q} .

4.1 EXACT BILEVEL FORMULATION

This subsection focuses on directly solving the bilevel optimization problem (4.1) to identify the optimal training distribution. Bilevel optimization has been extensively studied in the literature; see,

for example, the work of Colson et al. (2007); Golub & Pereyra (2003); Van Leeuwen & Aravkin (2021). For the infinite-dimensional problem (4.1), our approach begins by specifying a Hilbert space \mathcal{H} . Using its linear structure, we calculate the gradient of the Lagrangian associated with (4.1) with respect to the variable ν . This gradient computation leverages the *adjoint state method* and the envelope theorem (Afriat, 1971). In what follows, we specialize this general framework to the case where \mathcal{H} is a reproducing kernel Hilbert space (RKHS) with scalar output space $\mathcal{Y} = \mathbb{R}$. The resulting algorithm can be efficiently implemented using representer theorems for kernel methods (Micchelli & Pontil, 2005). Additionally, other linear classes of functions can be incorporated into this framework with minimal effort. Extending the framework to nonlinear, vector-valued function spaces \mathcal{H} , such as neural network or neural operator spaces, is an exciting direction for future research.

RKHS. Let $\mathcal{H} \coloneqq (\mathcal{H}_{\kappa}, \langle \cdot, \cdot \rangle_{\kappa}, \| \cdot \|_{\kappa})$ be a separable RKHS of real-valued functions on \mathcal{U} with bounded continuous reproducing kernel $\kappa \colon \mathcal{U} \times \mathcal{U} \to \mathbb{R}$. For fixed $\nu \in \mathscr{P}_2(\mathcal{U})$, let $\iota_{\nu} \colon \mathcal{H} \hookrightarrow L^2_{\nu}$ be the canonical inclusion map (which is compact under the assumptions on κ). The adjoint of the inclusion map is the kernel integral operator $\iota_{\nu}^* \colon L^2_{\nu} \to \mathcal{H}$ defined by $\iota_{\nu}^* h = \int_{\mathcal{U}} \kappa(\cdot\,,u) h(u) \nu(du)$. We also define the self-adjoint operator $\mathcal{K}_{\nu} \coloneqq \iota_{\nu}^* \iota_{\nu} \colon \mathcal{H} \to \mathcal{H}$, which has the same action as ι_{ν}^* .

Method of adjoints. Assume that $\mathcal{G}^* \in L^2_{\nu} \equiv L^2_{\nu}(\mathcal{U}; \mathbb{R})$ for every $\nu \in \mathscr{P}_2(\mathcal{U})$ under consideration; this is true, for example, if \mathcal{G}^* is uniformly bounded on \mathcal{U} . Define the data misfit $\Psi \colon \mathcal{H} \to \mathbb{R}$ by $\Psi(\mathcal{G}) \coloneqq \frac{1}{2} \| \iota_{\nu} \mathcal{G} - \mathcal{G}^* \|_{L^2_{\nu}}^2$; this is an abstract way to write the convex inner objective in (4.1) (rescaled by 1/2). By Lem. D.1, the inner solution $\widehat{\mathcal{G}}^{(\nu)}$ solves the equation $\mathcal{K}_{\nu} \widehat{\mathcal{G}}^{(\nu)} = \iota_{\nu}^* \mathcal{G}^*$ in \mathcal{H} . This fact allows us to define the Lagrangian $\mathcal{L} \colon \mathscr{P}_2(H) \times \mathcal{H} \times \mathcal{H} \to \mathbb{R}$ for (4.1) by

$$\mathcal{L}(\nu, \mathcal{G}, \lambda) = \frac{1}{2} \mathbb{E}_{\nu' \sim \mathbb{Q}} \| \iota_{\nu'} \mathcal{G} - \mathcal{G}^{\star} \|_{L_{\nu'}^{2}}^{2} + \left\langle (\mathcal{K}_{\nu} \mathcal{G} - \iota_{\nu}^{\star} \mathcal{G}^{\star}), \lambda \right\rangle_{\kappa}. \tag{4.2}$$

Differentiating \mathcal{L} with respect to $\mathcal{G} \in \mathcal{H}$ leads to the adjoint equation.

Lemma 4.2 (adjoint state equation). At a critical point of the map $\mathcal{G} \mapsto \mathcal{L}(\nu, \mathcal{G}, \lambda)$ for fixed (ν, \mathcal{G}) , the adjoint state $\lambda = \lambda^{(\nu)}(\mathcal{G}) \in \mathcal{H}$ satisfies the compact operator equation

$$\mathcal{K}_{\nu}\lambda^{(\nu)}(\mathcal{G}) = \mathbb{E}_{\nu' \sim \mathbb{Q}} \,\mathbb{E}_{u' \sim \nu'} \big[\kappa(\cdot, u') \big(\mathcal{G}^{\star}(u') - \mathcal{G}(u') \big) \big] \,. \tag{4.3}$$

The preceding results enable us to differentiate \mathcal{L} with respect to ν along the constraints of (4.1).

Proposition 4.3 (derivative: infinite-dimensional case). *Define* $\mathcal{J}: \mathscr{P}_2(\mathcal{U}) \to \mathbb{R}$ *by* $\mathcal{J}(\nu) := \mathcal{L}(\nu, \widehat{\mathcal{G}}^{(\nu)}, \lambda^{(\nu)}(\widehat{\mathcal{G}}^{(\nu)}))$. *Then, in the sense of duality,* $(D\mathcal{J})(\nu) = (\widehat{\mathcal{G}}^{(\nu)} - \mathcal{G}^*)\lambda^{(\nu)}(\widehat{\mathcal{G}}^{(\nu)})$ *for each* ν .

SM D.2 provides the proofs of Lem. 4.2 and Prop. 4.3.

Remark 4.4 (Wasserstein gradient flow). A nonparametric approach to solve the bilevel problem (4.1) is to derive the Wasserstein gradient flow and its Lagrangian equivalence $\dot{u} = -\nabla(D\mathcal{J})(\mathrm{Law}(u))(u)$, where $D\mathcal{J}$ is as in Prop. 4.3. This leads to a curve of measures $t \mapsto \mathrm{Law}(u(t))$ converging to a steady-state probability distribution. It is important to note that this dynamic requires gradients of the true map \mathcal{G}^* , which may not be available or could be too expensive to compute even if available. To implement this mean-field evolution, we can discretize it with interacting particles; cf. **SM** F.5.

Parametrized training distributions. Although Prop. 4.3 provides the desired derivative of the bilevel problem (4.1), it does not specify how to use the derivative to compute the optimal ν . This issue is nontrivial because $\nu \in \mathscr{P}_2(\mathcal{U})$ belongs to a nonlinear metric space. One approach assumes a parametrization $\nu = \nu_{\vartheta}$, where $\vartheta \in \mathcal{V}$ for some separable Hilbert parameter space \mathcal{V} . Typically $\mathcal{V} \subseteq \mathbb{R}^P$ for some $P \in \mathbb{N}$. We then use the chain rule to apply Euclidean gradient-based optimization tools to solve (4.1). To this end, define $J \colon \mathcal{V} \to \mathbb{R}$ by $\vartheta \mapsto J(\vartheta) \coloneqq \mathcal{J}(\nu_{\vartheta})$, where \mathcal{J} is as in Prop. 4.3. The following theorem computes the gradient $\nabla J \colon \mathcal{V} \to \mathcal{V}$.

Theorem 4.5 (gradient: parametric case). Suppose there exists a dominating measure μ_0 such that for every $\vartheta \in \mathcal{V}$, ν_{ϑ} has density p_{ϑ} with respect to μ_0 . If $(u,\vartheta) \mapsto p_{\vartheta}(u)$ is sufficiently regular, then

$$\nabla \mathsf{J}(\vartheta) = \int_{\mathcal{U}} (\widehat{\mathcal{G}}^{(\nu_{\vartheta})}(u) - \mathcal{G}^{\star}(u)) (\lambda^{(\nu_{\vartheta})}(\widehat{\mathcal{G}}^{(\nu_{\vartheta})}))(u) (\nabla_{\vartheta} \log p_{\vartheta}(u)) \nu_{\vartheta}(du). \tag{4.4}$$

In Thm. 4.5, μ_0 is typically the Lebesgue measure if $\mathcal{U} \subseteq \mathbb{R}^d$. SM D.2 contains the proof.

Example 4.6 (Gaussian parametrization). Consider the Gaussian model $\nu_{\vartheta} := \mathcal{N}(m_{\vartheta}, \mathcal{C}_{\vartheta})$ with μ_0 being the Lebesgue measure on $\mathcal{U} = \mathbb{R}^d$. Explicit calculations with Gaussian densities show that

$$\nabla_{\vartheta} \log p_{\vartheta}(u) = -\nabla_{\vartheta} \Phi(u, \vartheta), \text{ where } \Phi(u, \vartheta) := \frac{1}{2} \left\| \mathcal{C}_{\vartheta}^{-1/2}(u - m_{\vartheta}) \right\|_{2}^{2} + \frac{1}{2} \log \det(\mathcal{C}_{\vartheta}).$$
 (4.5)

Algorithm 1 Training Data Design via Gradient Descent on a Parametrized Bilevel Objective

- 1: **Initialize:** Parameter $\vartheta^{(0)}$, step sizes $\{t_k\}$
- 2: **for** $k = 0, 1, 2, \dots$ **do**

3: **Gradient step:** With ∇J as in (4.4), update the training distribution's parameters via

$$\vartheta^{(k+1)} = \vartheta^{(k)} - t_k \nabla \mathsf{J}(\vartheta^{(k)})$$

- 4: **if** stopping criterion is met **then**
- 5: Return $\vartheta^{(k+1)}$ then **break**

The parameters ϑ of ν_{ϑ} can now be updated using any gradient-based optimization method. For instance, a simple gradient flow in Hilbert space is given by $\dot{\vartheta} = -\nabla J(\vartheta)$.

Implementation. A gradient descent scheme for the bilevel optimization problem (4.1) is summarized in Alg. 1. Alternative gradient-based optimization methods can also be explored. Discretization details for the gradient (4.4) are provided in **SM** D.2; in particular, we "empiricalize" all expectations.

4.2 UPPER-BOUND MINIMIZATION

Although Alg. 1 is implementable assuming query access to the ground truth map $u\mapsto \mathcal{G}^\star(u)$, it is so far restricted to linear vector spaces \mathcal{H} . This subsection relaxes problem (4.1) by instead minimizing an upper bound on the objective, which is no longer a bilevel optimization problem. We then design an alternating descent algorithm that applies to any hypothesis class \mathcal{H} (e.g., neural networks and neural operators). Empirically, the alternating algorithm converges fast, often in a few iterations.

An alternating scheme. Recall the upper bound (3.4) from Cor. 3.2, which holds for any $\mu \in \mathscr{P}_2(\mathcal{U})$. The particular choice $\mu := \nu$, where ν is a candidate training data distribution, delivers a valid upper bound for the objective functional $\mathscr{E}_{\mathbb{Q}}(\widehat{\mathcal{G}}^{(\nu)})$ in (4.1). That is,

$$\mathscr{E}_{\mathbb{Q}}(\widehat{\mathcal{G}}^{(\nu)}) \leq \mathbb{E}_{u \sim \nu} \|\mathcal{G}^{\star}(u) - \widehat{\mathcal{G}}^{(\nu)}(u)\|_{\mathcal{Y}}^{2} + \sqrt{\mathbb{E}_{\nu' \sim \mathbb{Q}} c^{2}(\mathcal{G}^{\star}, \widehat{\mathcal{G}}^{(\nu)}, \nu, \nu')} \sqrt{\mathbb{E}_{\nu' \sim \mathbb{Q}} \mathsf{W}_{2}^{2}(\nu, \nu')}. \tag{4.6}$$

It is thus natural to approximate the solution to (4.1) by the minimizing distribution of the right-hand side of (4.6). The dependence of the factor c in Prop. 3.1 on the Lipschitz constant of $\widehat{\mathcal{G}}^{(\nu)}$ is hard to capture in practice. One way to circumvent this difficulty is to further assume an upper bound $\operatorname{Lip}(\widehat{\mathcal{G}}^{(\nu)}) \leq R$ uniformly in ν for some constant R>0. We define $c_R(\mathcal{G}^\star,\nu,\nu')$ to be $c(\mathcal{G}^\star,\widehat{\mathcal{G}}^{(\nu)},\nu,\nu')$ with all instances of $\operatorname{Lip}(\widehat{\mathcal{G}}^{(\nu)})$ replaced by R. The assumption of a uniform bound on the Lipschitz constant is valid for certain chosen model classes \mathcal{H} , e.g., kernels with Lipschitz feature maps or Lipschitz-constrained neural networks (Gouk et al., 2021).

We denote the candidate set of training distributions by $P \subseteq \mathscr{P}_2(\mathcal{U})$. Finally, the upper-bound minimization problem becomes

$$\inf_{(\nu,\mathcal{G})\in\mathsf{P}\times\mathcal{H}} \left\{ \mathbb{E}_{u\sim\nu} \left\| \mathcal{G}^{\star}(u) - \mathcal{G}(u) \right\|_{\mathcal{Y}}^{2} + \sqrt{\mathbb{E}_{\nu'\sim\mathbb{Q}} c_{R}^{2} (\mathcal{G}^{\star},\nu,\nu')} \sqrt{\mathbb{E}_{\nu'\sim\mathbb{Q}} \mathsf{W}_{2}^{2}(\nu,\nu')} \right\}. \tag{4.7}$$

Note that since only the first term of (4.7) depends on the variable \mathcal{G} , the minimizing \mathcal{G} is given by $\widehat{\mathcal{G}}^{(\nu)}$, as defined previously. This observation motivates the alternating scheme outlined in Alg. 2.

It is important to note that not all alternating algorithms have convergence guarantees, which requires a detailed understanding of the associated fixed-point operator. In contrast, the proposed Alg. 2 comes with the following convergence guarantee.

Proposition 4.7 (monotonicity). *Under Alg. 2, the objective in* (4.7) *is always non-increasing.*

Proof. Line 3 ensures that the first term of the objective (4.7) does not increase, while keeping the second term constant. Next, Line 4 minimizes both terms. Iterating this argument, we deduce that the algorithm maintains a non-increasing objective value for all k.

Remark 4.8 (connection to Wasserstein barycenters). If the set P contains distributions with uniformly bounded second moments and the first term on the right-hand side of (4.6) is uniformly bounded over P, then the optimal value of (4.7) is bounded above by the value of another minimization problem

Algorithm 2 Training Data Design via Alternating Model Fitting and Distribution Update Steps

- 1: **Initialize:** Training data distribution $\nu^{(0)}$, estimate of the constant R > 0
- 2: **for** $k = 0, 1, 2, \dots$ **do**

3: **Model training:** To obtain the trained model $\widehat{\mathcal{G}}^{(k)}$, solve the optimization problem

$$\widehat{\mathcal{G}}^{(k)} \leftarrow \operatorname*{arg\,min}_{\mathcal{G} \in \mathcal{H}} \mathbb{E}_{u \sim \nu^{(k)}} \left\| \mathcal{G}^{\star}(u) - \mathcal{G}(u) \right\|_{\mathcal{Y}}^{2}$$

4: **Distribution update:** Update the training distribution by solving

$$\nu^{(k+1)} \leftarrow \mathop{\arg\min}_{\nu \in \mathsf{P}} \Bigl\{ \mathbb{E}_{u \sim \nu} \bigl\| \mathcal{G}^{\star}(u) - \widehat{\mathcal{G}}^{(k)}(u) \bigr\|_{\mathcal{Y}}^2 + \sqrt{\mathbb{E}_{\nu' \sim \mathbb{Q}} \, c_R^2 \bigl(\mathcal{G}^{\star}, \nu, \nu' \bigr)} \sqrt{\mathbb{E}_{\nu' \sim \mathbb{Q}} \, \mathsf{W}_2^2(\nu, \nu')} \Bigr\}$$

- 5: **if** stopping criterion is met **then**
- 6: Return $(\widehat{\mathcal{G}}^{(k)}, \nu^{(k+1)})$ then **break**

whose solution in $\mathscr{P}_2(\mathcal{U})$ is the 2-Wasserstein barycenter of $\mathbb{Q} \in \mathscr{P}_2(\mathscr{P}_2(\mathcal{U}))$ (Backhoff-Veraguas et al., 2022, Thm 3.7, p. 444). An analogous bound is even easier to verify for the W_1 result (3.1). Thus, the \mathcal{G}^* -dependent solution of the proposed problem (4.7) over P (and hence also the solution of Eqn. (4.1)) is always *at least as good* as the barycenter, which only depends on P and \mathbb{Q} .

Implementation. To implement Alg. 2, we simply replace expectations with finite averages over empirical samples. The number of samples per iteration is a hyperparameter. Line 3 in Alg. 2 is the usual model training step, i.e., Eqn. (2.1), while the distribution update step in Line 4 can be achieved in several ways. One way involves interacting particle systems derived from (Wasserstein) gradient flows over $P \subseteq \mathscr{P}_2(\mathcal{U})$. Alternatively, if the distribution is parametrized, we can optimize over the parametrization, e.g., transport maps or Gaussian means and covariances. In this case, one can apply global optimization methods or Euclidean gradient-based solvers.

To apply gradient-based optimization algorithms, we need the derivative of the objective in (4.7) with respect to ν . Up to constants, the objective takes the form $F \colon P \subseteq \mathscr{P}_2(\mathcal{U}) \to \mathbb{R}$, where

$$\nu \mapsto F(\nu) := \mathbb{E}_{u \sim \nu} f(u) + \sqrt{1 + \mathsf{m}_2(\nu)} \sqrt{\mathbb{E}_{\nu' \sim \mathbb{Q}} \mathsf{W}_2^2(\nu, \nu')}$$
(4.8)

and $f: \mathcal{U} \to \mathbb{R}$. The next result differentiates F; the proof is in SM D.3.

Proposition 4.9 (derivative: infinite-dimensional case). Let F be as in (4.8) with f and P sufficiently regular. Then for every $\nu \in P$ and $u \in U$, it holds that

$$(DF(\nu))(u) = f(u) + \frac{1}{2} \sqrt{\frac{\mathbb{E}_{\nu' \sim \mathbb{Q}} \, \mathsf{W}_2^2(\nu, \nu')}{1 + \mathsf{m}_2(\nu)}} \|u\|_{\mathcal{U}}^2 + \sqrt{\frac{1 + \mathsf{m}_2(\nu)}{\mathbb{E}_{\nu' \sim \mathbb{Q}} \, \mathsf{W}_2^2(\nu, \nu')}} \, \mathbb{E}_{\nu' \sim \mathbb{Q}} \, \phi_{\nu, \nu'}(u)$$
 (4.9)

in the sense of duality. The potential $\phi_{\nu,\nu'} \colon \mathcal{U} \to \mathbb{R}$ satisfies $\phi_{\nu,\nu'} = \frac{1}{2} \| \cdot \|_{\mathcal{U}}^2 - \varphi_{\nu,\nu'}$, where $\nabla \varphi_{\nu,\nu'} = T_{\nu \to \nu'}$ and $T_{\nu \to \nu'}$ is the W_2 -optimal transport map from ν to ν' .

Taking the gradient of $u \mapsto DF(\nu)(u)$ in (4.9) yields the Wasserstein gradient, which can be utilized in Wasserstein gradient flow discretizations; see Rmk. 4.4. In the case of parametric distributions, combining Prop. 4.9 with the chain rule yields gradient descent schemes similar to those introduced in Sec. 4.1. For further details, see Lem. D.8 and its special case derived from Rmk. D.9 in **SM** D.3.

5 Numerical results

This section implements Alg. 1 (the bilevel gradient descent algorithm) for function approximation using kernel methods and Alg. 2 (the alternating minimization algorithm) for two operator learning tasks. Comprehensive experiment details and settings are collected in **SM** F.

Bilevel gradient descent algorithm. We instantiate a discretization of Alg. 1 on several function approximation benchmarks. The test functions $\mathcal{G}^* \in \{g_i\}_{i=1}^4$ each map \mathbb{R}^d to \mathbb{R} and are comprised of the Sobol G (g_1) and Friedmann functions (g_2, g_3) , and a kernel expansion (g_4) . We use the kernel framework from Sec. 4.1 with \mathcal{H} the RKHS of the squared exponential kernel with a fixed scalar bandwidth hyperparameter. Also, \mathbb{Q} is an empirical measure with 10 atoms centered at fixed Gaussian

Table 1: Function approximation with bilevel Alg. 1. Reported values are mean Err for kernel regressors trained with 1024 labeled data samples from the optimal ν_{ϑ} from Alg. 1 (1000 iterations) vs. nonadaptive sampling distributions. Two standard deviations from the mean Err over 10 independent runs is also reported.

	$g_1 \ (d=2)$	$g_2 \ (d=5)$	$g_2\ (d=8)$	$g_3 (d=4)$	$g_3 \ (d=5)$	$g_4 \ (d=10)$
Optimized	0.039 ± 0.007	0.189 ± 0.004	0.556 ± 0.007	0.271 ± 0.017	0.557 ± 0.006	0.020 ± 0.001
Normal	0.809 ± 0.018	0.812 ± 0.009	0.949 ± 0.003	0.939 ± 0.006	0.949 ± 0.003	0.030 ± 0.001
Barycenter	0.541 ± 0.018	0.518 ± 0.014	0.782 ± 0.008	0.769 ± 0.009	0.784 ± 0.007	0.021 ± 0.001
Mixture	0.308 ± 0.011	0.372 ± 0.015	0.712 ± 0.011	0.471 ± 0.037	0.658 ± 0.011	0.022 ± 0.001
Uniform	0.983 ± 0.001	0.971 ± 0.002	0.993 ± 0.000	0.998 ± 0.000	1.000 ± 0.000	0.121 ± 0.001

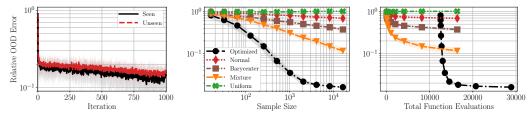


Figure 2: Alg. 1 applied to ground truth $g_1 \colon \mathbb{R}^2 \to \mathbb{R}$. (Left) Evolution of Err over 1000 iterations (N=250 samples per step) of gradient descent. (Center) Err of model trained on N samples from optimized ν_{ϑ} (ours) vs. initial normal $\mathcal{N}(m_0, I_2)$, empirical \mathbb{Q} W₂-barycenter, empirical \mathbb{Q} mixture, and $\mathrm{Unif}([0,1]^2)$. (Right) Same as center, except incorporating the additional function evaluation cost incurred from Alg. 1. Shading represents two standard deviations away from the mean Err over 10 independent runs.

measures realized by standard normally distributed means and Wishart distributed covariance matrices. A cosine annealing scheduler determines the gradient descent step sizes. The training distribution $\nu_{\vartheta} := \mathcal{N}(m, \mathcal{C})$ is parametrized by the pair $\vartheta = (m, \mathcal{C})$, which we optimize using Alg. 1.

Tab. 1 shows that models trained on the final iterate ν_{ϑ} of Alg. 1 always outperforms those trained on nonadaptive distributions in terms of the root relative average OOD squared error $\text{Err} := (\mathscr{E}_{\mathbb{Q}}(\widehat{\mathcal{G}}^{(\nu)})/\mathscr{E}_{\mathbb{Q}}(0))^{1/2}$. However, the performance gap can close depending on the complexity of \mathbb{Q} or \mathcal{G}^{\star} . Fig. 2 specializes to the $\mathcal{G}^{\star} = g_1$ case. Although fitting the model on data sampled from the precomputed optimal ν_{ϑ} uniformly outperforms nonadaptive distribution sampling, when corrected for the "offline" cost of finding the optimal ν_{ϑ} , the gains are more modest (Fig. 2, right).

Alternating minimization algorithm. In the first two examples, we apply the alternating scheme Alg. 2 to two important operator learning problems in the computational sciences. Let $\Omega \subset \mathbb{R}^2$ denote a bounded domain with boundary $\partial\Omega$. Consider the elliptic PDE $\nabla \cdot (a(x)\nabla u(x))=0$, where $x\in\Omega$, a=a(x)>0 is the conductivity, and u=u(x) is the PDE solution. Given a mean-zero Neumann boundary condition (BC) $g=a\frac{\partial u}{\partial \mathbf{n}}|_{\partial\Omega}$, one can obtain a unique PDE solution u and the resulting Dirichlet data $f:=u|_{\partial\Omega}$. For a fixed conductivity a, we define a Neumann-to-Dirichlet (NtD) map $\Lambda_a^{-1}:g\mapsto f$, which is also used in the EIT problem (Dunlop & Stuart, 2016); see also Sec. 1. On the other hand, the map $\mathcal{G}:a\mapsto u$ with a fixed Neumann or Dirichlet BC (and nonzero source) is the Darcy flow parameter-to-solution operator, which plays a crucial role in geothermal energy extraction and carbon storage applications. Another important operator maps the scattering coefficient to the spatial-domain density in the radiative transport equation, which has key applications in optical tomography and atmospheric science. Learning surrogate operators for these mappings can enable rapid simulations for efficient inference and control.

NtD map learning. In this example, we aim to find the optimal training distribution for learning the NtD map using the alternating minimization algorithm (AMA). The domain $\Omega=(0,1)^2$ and the conductivity a is fixed. The meta test distribution is a mixture of K components $\mathbb{Q}=\frac{1}{K}\sum_{k=1}^K \delta_{\nu_k'}$, where each ν_k' corresponds to a distribution over g, the Neumann BC. Fig. 3 illustrates the reduction in both average relative OOD error (3.3) and relative AMA loss (4.7) over 10 iterations of Alg. 2.

Darcy flow forward map learning. We use AMA to identify the optimal training distribution for learning the Darcy flow parameter-to-solution map. Here, $\mathbb Q$ is as before except now each ν_k' represents a test distribution over the conductivity. Fig. 4 illustrates improved predictions of the neural operator trained on distributions obtained from Alg. 2 across various iterations.

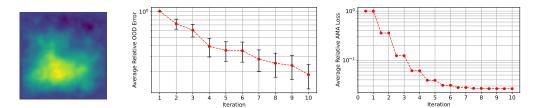


Figure 3: (Left) Conductivity for NtD learning. (Center) After 80 independent runs, decay of the average relative OOD error of the model when trained on the optimal distribution identified at each iteration of Alg. 2; a 95% confidence interval of the true relative OOD error is provided at each iteration. (Right) Decay of average AMA loss defined in (4.7) vs. iteration relative to the same loss at initialization.

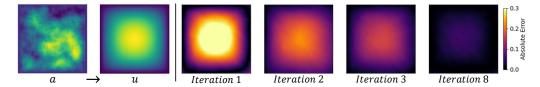


Figure 4: The first two images show the test conductivity a and true PDE solution u. The next four images show the absolute error $|u_{pred} - u|$ of the model trained on the training distribution from iterations 1, 2, 3, and 8.

Radiative transport operator learning. We apply a nonparametric particle-based version of AMA to find the optimal training distribution for learning the solution operator of the radiative transport equation. We vary the Knudsen number which corresponds to different physical regimes. Fig. 5 shows how the model trained on the optimal empirical measure compares to the initial model and a benchmark model trained on the empirical $\mathbb Q$ mixture. When the sample size is N=336 particles, a single AMA iteration reduces the average relative OOD error by 85% compared to its initial value.

6 Conclusion

This paper shows how optimizing training distributions using theoretically founded bilevel and alternating optimization algorithms produces more accurate and robust models for function approximation and PDE operator learning. Numerical experiments instate the methods with parametric distribution families and nonparametric particle-based discretizations of Wasserstein gradient flows. Though promising, the proposed approach is limited by distribution family expressivity, global optimization, and line search steps. Future work will address these limitations by incorporating transport maps and importance reweighting techniques. Understanding cost versus accuracy trade-offs for sample allocation and iteration complexity of the methodology is another important future challenge.

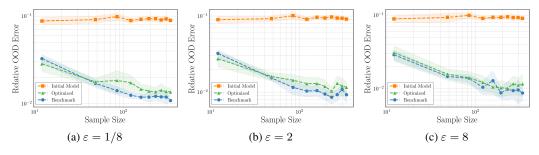


Figure 5: Relative OOD error vs. sample size N for learning the radiative transport solution operator. Results are shown for Knudsen numbers $\varepsilon \in \{1/8, 2, 8\}$. For each N, each panel displays 95% confidence intervals over 10 trials for three DeepONet models trained for 5000 epochs: the initial model, the model after one particle-based AMA iteration, and a benchmark trained the test distribution mixture $\nu_{\mathbb{Q}} \coloneqq \frac{1}{3} \sum_{k=1}^{3} \nu_{k}'$.

REPRODUCIBILITY STATEMENT

We have taken several steps to ensure the reproducibility of our results. All proposed algorithms are described in detail in the main text, with pseudocode and additional implementation and discretization details provided in **SM** D, E, and F. Complete proofs of all theoretical results from the main text, together with explicit statements of assumptions, interpretations, and discussion, are included in **SM** C and D. The function approximation and PDE benchmarks are introduced in the main text, with full specifications of parameter ranges, boundary conditions, and discretizations in **SM** E and F. An anonymous numerical implementation, including code for training distribution optimization, model training, and evaluation, is provided in the supplementary material as a zip file. Hyperparameters and experimental settings are listed in **SM** F and in configuration files within the code repository. Together, these resources should enable full replication and extension of our results.

REFERENCES

- Ben Adcock. Optimal sampling for least-squares approximation. preprint arXiv:2409.02342, 2024.
- Ben Adcock, Juan M Cardenas, and Nick Dexter. CAS4DL: Christoffel adaptive sampling for function approximation via deep learning. *Sampling Theory, Signal Processing, and Data Analysis*, 20(2):21, 2022.
- SN Afriat. Theory of maxima and the method of Lagrange. *SIAM Journal on Applied Mathematics*, 20(3):343–357, 1971.
- Sergios Agapiou, Omiros Papaspiliopoulos, Daniel Sanz-Alonso, and Andrew M Stuart. Importance sampling: Intrinsic dimension and computational cost. *Statistical Science*, pp. 405–431, 2017.
- Pedro C Álvarez-Esteban, E Del Barrio, JA Cuesta-Albertos, and C Matrán. A fixed-point approach to barycenters in Wasserstein space. *Journal of Mathematical Analysis and Applications*, 441(2): 744–762, 2016.
- Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *preprint arXiv:1907.02893*, 2019.
- Anthony Atkinson, Alexander Donev, and Randall Tobias. *Optimum experimental designs, with SAS*, volume 34. OUP Oxford, 2007.
- Julio Backhoff-Veraguas, Joaquin Fontbona, Gonzalo Rios, and Felipe Tobar. Bayesian learning with Wasserstein barycenters. *ESAIM: Probability and Statistics*, 26:436–472, 2022.
- Andrew R Barron, Albert Cohen, Wolfgang Dahmen, and Ronald A DeVore. Approximation and learning by greedy algorithms. *The Annals of Statistics*, 36(1), 2008.
- Jose Antonio Lara Benitez, Takashi Furuya, Florian Faucher, Anastasis Kratsios, Xavier Tricoche, and Maarten V de Hoop. Out-of-distributional risk bounds for neural operators with applications to the Helmholtz equation. *Journal of Computational Physics*, art. 113168, 2024.
- Nicolas Boullé and Alex Townsend. A mathematical guide to operator learning. In *Handbook of Numerical Analysis*, volume 25, pp. 83–125. Elsevier, 2024.
- Nicolas Boullé, Diana Halikias, and Alex Townsend. Elliptic PDE learning is provably data-efficient. *Proceedings of the National Academy of Sciences*, 120(39):e2303904120, 2023.
- Russel E Caflisch. Monte Carlo and quasi-Monte Carlo methods. *Acta Numerica*, 7:1–49, 1998.
- Wenqian Chen, Amanda A Howard, and Panos Stinis. Self-adaptive weights based on balanced residual decay rate for physics-informed neural networks and deep operator networks. *Journal of Computational Physics*, art. 114226, 2025a.
- Yifan Chen, Ethan N Epperly, Joel A Tropp, and Robert J Webber. Randomly pivoted Cholesky: Practical approximation of a kernel matrix with few entry evaluations. *Communications on Pure and Applied Mathematics*, 78(5):995–1041, 2025b.

Sinho Chewi. *Log-Concave Sampling*. Book Draft, 2024.

- Jonathan Cockayne and Andrew Duncan. Probabilistic gradients for fast calibration of differential equation models. *SIAM/ASA Journal on Uncertainty Quantification*, 9(4):1643–1672, 2021.
 - Albert Cohen and Giovanni Migliorati. Optimal weighted least-squares methods. *The SMAI Journal of Computational Mathematics*, 3:181–203, 2017.
 - Benoît Colson, Patrice Marcotte, and Gilles Savard. An overview of bilevel optimization. *Annals of Operations Research*, 153:235–256, 2007.
 - Masoumeh Dashti and Andrew M Stuart. The Bayesian Approach to Inverse Problems. *Handbook of Uncertainty Quantification*, pp. 311–428, 2017.
 - Maarten V de Hoop, Nikola B Kovachki, Nicholas H Nelsen, and Andrew M Stuart. Convergence rates for learning linear operators from noisy data. *SIAM/ASA Journal on Uncertainty Quantification*, 11(2):480–513, 2023.
 - Songgaojun Deng, Olivier Sprangers, Ming Li, Sebastian Schelter, and Maarten De Rijke. Domain generalization in time series forecasting. *ACM Transactions on Knowledge Discovery from Data*, 18(5):1–24, 2024.
 - Xun Deng, Wenjie Wang, Fuli Feng, Hanwang Zhang, Xiangnan He, and Yong Liao. Counterfactual active learning for out-of-distribution generalization. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 11362–11377, 2023.
 - Josef Dick, Frances Y Kuo, and Ian H Sloan. High-dimensional integration: The quasi-Monte Carlo way. *Acta Numerica*, 22:133–288, 2013.
 - Oliver RA Dunbar, Nicholas H Nelsen, and Maya Mutic. Hyperparameter optimization for randomized algorithms: a case study on random features. *Statistics and Computing*, 35(3):1–28, 2025.
 - Matthew M Dunlop and Andrew M Stuart. The Bayesian formulation of EIT: Analysis and algorithms. *Inverse Problems and Imaging*, 10(4):1007–1036, 2016.
 - Cian Eastwood, Alexander Robey, Shashank Singh, Julius Von Kügelgen, Hamed Hassani, George J Pappas, and Bernhard Schölkopf. Probable domain generalization via quantile risk minimization. *Advances in Neural Information Processing Systems*, 35:17340–17358, 2022.
 - Ethan Epperly and Elvira Moreno. Kernel quadrature with randomly pivoted Cholesky. In *Advances in Neural Information Processing Systems*, volume 36, pp. 65850–65868, 2023.
 - Yuwei Fan and Lexing Ying. Solving electrical impedance tomography with deep learning. *Journal of Computational Physics*, 404:109119, 2020.
 - Tongtong Fang, Nan Lu, Gang Niu, and Masashi Sugiyama. Generalizing importance weighting to a universal solver for distribution shift problems. *Advances in Neural Information Processing Systems*, 36:24171–24190, 2023.
 - Valerii Fedorov. Optimal experimental design. Wiley Interdisciplinary Reviews: Computational Statistics, 2(5):581–589, 2010.
 - Adam Foster, Martin Jankowiak, Elias Bingham, Paul Horsfall, Yee Whye Teh, Thomas Rainforth, and Noah Goodman. Variational Bayesian optimal experimental design. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- Jean-Christophe Gagnon-Audet, Kartik Ahuja, Mohammad-Javad Darvishi Bayazi, Pooneh Mousavi,
 Guillaume Dumas, and Irina Rish. WOODS: Benchmarks for Out-of-Distribution Generalization
 in Time Series. *Trans. Mach. Learn. Res.*, 2023.
 - Zhiwei Gao, Liang Yan, and Tao Zhou. Failure-informed adaptive sampling for PINNs. *SIAM Journal on Scientific Computing*, 45(4):A1971–A1994, 2023.

- Alfredo Garbuno-Inigo, Tapio Helin, Franca Hoffmann, and Bamdad Hosseini. Bayesian posterior perturbation analysis with integral probability metrics. *preprint arXiv:2303.01512*, 2023.
 - Matthias Gelbrich. On a formula for the L^2 Wasserstein metric between measures on Euclidean and Hilbert spaces. *Mathematische Nachrichten*, 147(1):185–203, 1990.
 - Gene Golub and Victor Pereyra. Separable nonlinear least squares: the variable projection method and its applications. *Inverse problems*, 19(2):R1, 2003.
 - Henry Gouk, Eibe Frank, Bernhard Pfahringer, and Michael J Cree. Regularisation of neural networks by enforcing Lipschitz continuity. *Machine Learning*, 110:393–416, 2021.
 - Philipp Grohs, Samuel Lanthaler, and Margaret Trautner. Theory-to-practice gap for neural networks and neural operators. *preprint arXiv:2503.18219*, 2025.
 - Timothy Hospedales, Antreas Antoniou, Paul Micaelli, and Amos Storkey. Meta-learning in neural networks: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 44(9): 5149–5169, 2021.
 - Xun Huan, Jayanth Jagalur, and Youssef Marzouk. Optimal experimental design: Formulations and computations. *Acta Numerica*, 33:715–840, 2024.
 - Ferenc Huszár and David Duvenaud. Optimally-weighted herding is Bayesian quadrature. In *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence*, pp. 377–386, 2012.
 - Ruhui Jin, Martin Guerra, Qin Li, and Stephen Wright. Optimal design for linear models via gradient flow. *preprint arXiv:2401.07806*, 2024a.
 - Ruhui Jin, Qin Li, Stephen O Mussmann, and Stephen J Wright. Continuous nonlinear adaptive experimental design with gradient flow. *preprint arXiv:2411.14332*, 2024b.
 - Olav Kallenberg. Random measures, theory and applications, volume 1. Springer, 2017.
 - Pritish Kamath, Akilesh Tangella, Danica Sutherland, and Nathan Srebro. Does invariant risk minimization capture invariance? In *International Conference on Artificial Intelligence and Statistics*, pp. 4069–4077. PMLR, 2021.
 - Nikola B Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew M Stuart, and Anima Anandkumar. Neural operator: Learning maps between function spaces with applications to PDEs. *Journal of Machine Learning Research*, 24(89):1–97, 2023.
 - Nikola B Kovachki, Samuel Lanthaler, and Hrushikesh Mhaskar. Data complexity estimates for operator learning. *preprint arXiv:2405.15992*, 2024a.
 - Nikola B Kovachki, Samuel Lanthaler, and Andrew M Stuart. Operator learning: Algorithms and analysis. *Handbook of Numerical Analysis*, 25:419–467, 2024b.
 - David Krieg, Erich Novak, and Mathias Sonnleitner. Recovery of Sobolev functions restricted to iid sampling. *Mathematics of Computation*, 91(338):2715–2738, 2022.
 - Shibo Li, Xin Yu, Wei Xing, Robert Kirby, Akil Narayan, and Shandian Zhe. Multi-resolution active learning of Fourier neural operators. In *International Conference on Artificial Intelligence and Statistics*, pp. 2440–2448. PMLR, 2024.
 - Zongyi Li, Nikola B. Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew M. Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. In *International Conference on Learning Representations*, 2021.
 - Fengming Lin, Xiaolei Fang, and Zheming Gao. Distributionally robust optimization: A review on theory and applications. *Numerical Algebra, Control and Optimization*, 12(1):159–212, 2022a.
 - Yong Lin, Hanze Dong, Hao Wang, and Tong Zhang. Bayesian invariant risk minimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 16021–16030, 2022b.

- Ilya Loshchilov and Frank Hutter. SGDR: Stochastic Gradient Descent with Warm Restarts. In *International Conference on Learning Representations*, 2017.
 - Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3(3):218–229, 2021.
- Kjetil O Lye, Siddhartha Mishra, Deep Ray, and Praveen Chandrashekar. Iterative surrogate model optimization (ISMO): An active learning algorithm for PDE constrained optimization with deep neural networks. *Computer Methods in Applied Mechanics and Engineering*, 374:113575, 2021.
- Neil Rohit Mallinar, Austin Zane, Spencer Frei, and Bin Yu. Minimum-norm interpolation under covariate shift. In *International Conference on Machine Learning*, pp. 34543–34585. PMLR, 2024.
- Levi D McClenny and Ulisses M Braga-Neto. Self-adaptive physics-informed neural networks. *Journal of Computational Physics*, 474:111722, 2023.
- Charles A Micchelli and Massimiliano Pontil. On learning vector-valued functions. *Neural Computation*, 17(1):177–204, 2005.
- Roberto Molinaro, Yunan Yang, Björn Engquist, and Siddhartha Mishra. Neural inverse operators for solving PDE inverse problems. In *International Conference on Machine Learning*, pp. 25105–25139. PMLR, 2023.
- Daniel Musekamp, Marimuthu Kalimuthu, David Holzmüller, Makoto Takamoto, and Mathias Niepert. Active learning for neural PDE solvers. *preprint arXiv:2408.01536*, 2024.
- Nicholas H Nelsen and Andrew M Stuart. Operator learning using random features: A tool for scientific computing. *SIAM Review*, 66(3):535–571, 2024.
- Nicholas H Nelsen and Yunan Yang. Operator learning meets inverse problems: A probabilistic perspective. *preprint arXiv:2508.20207*, 2025.
- Weihang Ouyang, Min Zhu, Wei Xiong, Si-Wei Liu, and Lu Lu. RAMS: Residual-based adversarial-gradient moving sample method for scientific machine learning in solving partial differential equations. *preprint arXiv:2509.01234*, 2025.
- Hyunkyu Park, Kyungseo Park, Sangwoo Mo, and Jung Kim. Deep neural network based electrical impedance tomographic sensing methodology for large-area robotic tactile sensing. *IEEE Transactions on Robotics*, 37(5):1570–1583, 2021.
- Ethan Pickering, Stephen Guth, George Em Karniadakis, and Themistoklis P Sapsis. Discovering and forecasting extreme events via active learning in neural operators. *Nature Computational Science*, 2(12):823–833, 2022.
- Daniel Potts and Michael Schmischke. Interpretable approximation of high-dimensional data. *SIAM Journal on Mathematics of Data Science*, 3(4):1301–1323, 2021.
- Hamed Rahimian and Sanjay Mehrotra. Frameworks and results in distributionally robust optimization. *Open Journal of Mathematical Optimization*, 3:1–85, 2022.
- Alessandro Rudi, Raffaello Camoriano, and Lorenzo Rosasco. Less is more: Nyström computational regularization. In *Advances in Neural Information Processing Systems*, volume 28, 2015.
- Alessandro Rudi, Daniele Calandriello, Luigi Carratino, and Lorenzo Rosasco. On fast leverage score sampling and optimal learning. In *Advances in Neural Information Processing Systems*, volume 31, 2018.
- Esha Saha, Hayden Schaeffer, and Giang Tran. HARFE: hard-ridge random feature expansion. Sampling Theory, Signal Processing, and Data Analysis, 21(2):27, 2023.
- Filippo Santambrogio. *Optimal Transport for Applied Mathematicians: Calculus of Variations, PDEs, and Modeling*, volume 87. Birkhäuser, 2015.

- Paul Setinek, Gianluca Galletti, Thomas Gross, Dominik Schnürer, Johannes Brandstetter, and
 Werner Zellinger. SIMSHIFT: A Benchmark for Adapting Neural Surrogates to Distribution Shifts.
 preprint arXiv:2506.12007, 2025.
 - Burr Settles. Active learning literature survey. Computer Sciences Technical Report 1648, University of Wisconsin–Madison, 2009.
 - Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P Adams, and Nando De Freitas. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1): 148–175, 2015.
 - Samuli Siltanen, Jennifer Mueller, and David Isaacson. An implementation of the reconstruction algorithm of A. Nachman for the 2D inverse conductivity problem. *Inverse Problems*, 16(3):681, 2000.
 - Anurag Singh, Siu Lun Chau, Shahine Bouabid, and Krikamol Muandet. Domain generalisation via imprecise learning. In *International Conference on Machine Learning*, pp. 45544–45570. PMLR, 2024.
 - Andrew M Stuart. Inverse problems: A Bayesian perspective. Acta Numerica, 19:451–559, 2010.
 - Unique Subedi and Ambuj Tewari. On the benefits of active data collection in operator learning. *preprint arXiv:2410.19725*, 2024.
 - Kejun Tang, Xiaoliang Wan, and Chao Yang. Das-pinns: A deep adaptive sampling method for solving high-dimensional partial differential equations. *Journal of Computational Physics*, 476: 111868, 2023.
 - Juan Diego Toscano, Daniel T Chen, Vivek Oommen, and George Em Karniadakis. A variational framework for residual-based adaptivity in neural PDE solvers and operator learning. *preprint arXiv:2509.14198*, 2025.
 - Tristan Van Leeuwen and Aleksandr Y Aravkin. Variable projection for nonsmooth problems. *SIAM Journal on Scientific Computing*, 43(5):S249–S268, 2021.
 - Cédric Villani. Topics in optimal transportation, volume 58. American Mathematical Society, 2021.
 - Jindong Wang, Cuiling Lan, Chang Liu, Yidong Ouyang, Tao Qin, Wang Lu, Yiqiang Chen, Wenjun Zeng, and Philip S Yu. Generalizing to unseen domains: A survey on domain generalization. *IEEE transactions on knowledge and data engineering*, 35(8):8052–8072, 2022a.
 - Sifan Wang, Hanwen Wang, and Paris Perdikaris. Improved architectures and training algorithms for deep operator networks. *Journal of Scientific Computing*, 92(2):35, 2022b.
 - Chenxi Wu, Min Zhu, Qinyang Tan, Yadhu Kartha, and Lu Lu. A comprehensive study of non-adaptive and residual-based adaptive sampling for physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*, 403:115671, 2023.
 - Yang Yang, Yuxuan Zhang, Xin Song, and Yi Xu. Not all out-of-distribution data are harmful to open-set active learning. In *Advances in Neural Information Processing Systems*, volume 36, pp. 13802–13818, 2023.
 - Xiao Zhou, Yong Lin, Weizhong Zhang, and Tong Zhang. Sparse invariant risk minimization. In *International Conference on Machine Learning*, pp. 27222–27244. PMLR, 2022.
 - Dixian Zhu, Zhe Li, Xiaoyu Wang, Boqing Gong, and Tianbao Yang. A robust zero-sum game framework for pool-based active learning. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 517–526. PMLR, 2019.

758

759 760

761 762

764

765

766

767

768

769

770

771

772

773

774

775

776

777

778

779

781

782

783

784

785

786

787

788

789

790

791

792

793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

Supplementary Material for:

Learning Where to Learn:
Training Data Distribution Optimization for Scientific Machine Learning

A EXPANDED OVERVIEW OF RELATED WORK

Beyond the SciML literature reviewed in Sec. 1, our work also sits alongside research on learning under distribution shift (Mallinar et al., 2024). This includes meta-learning, domain generalization and adaptation, and active learning and optimal experimental design (OED). We now explain these connections.

Meta-learning aims to learn models or priors that adapt rapidly across tasks drawn from a task distribution (Hospedales et al., 2021). While related in spirit, meta-learning typically treats the data-generating distribution per task as given; in contrast, we *optimize* the training distribution itself to improve average performance across deployment regimes. Both frameworks involve a two-level structure: an inner optimization that optimizes the model over a training dataset, and an outer optimization that adjusts a higher-level quantity beyond model parameters, guided by generalization performance. In optimization-based meta-learning methods, the outer loop adjusts the model initialization or update rules, allowing for rapid adaptation across new tasks. On the other hand, the outer loop in our work optimizes the training data distribution.

Domain generalization seeks representations that transfer across multiple source domains without access to target data (Wang et al., 2022a; Deng et al., 2024; Gagnon-Audet et al., 2023; Eastwood et al., 2022; Singh et al., 2024). Instead, domain adaptation leverages (often unlabeled) target data to reduce source-target discrepancy (Fang et al., 2023). Both paradigms assume a fixed training set and focus on achieving robustness during training. Common benchmarks for unsupervised domain adaptation only assume access to samples from a single target domain $(\nu'$ in $\mathbb{Q} = \delta_{\nu'})$ and do not have the ability to generate new labels (Setinek et al., 2025); this setting is orthogonal to our work. Robust domain generalization methods, such as distributionally robust optimization and invariant risk minimization (Lin et al., 2022a; Rahimian & Mehrotra, 2022; Arjovsky et al., 2019; Kamath et al., 2021; Lin et al., 2022b; Zhou et al., 2022), modify objectives or constraints to ensure stability across domains given the training data. In contrast, we design the training distribution before learning: we retain standard training objectives but actively shape where data are drawn for improved sample complexity and robustness to distribution shift in regression problems. We do so with bilevel (or alternating) optimization in the space of probability measures, guided by theoretical bounds. In particular, the minimizing training data distribution depends not just on the chosen model architecture and the accuracy metric, but also on the target map itself.

The perspective presented in our paper naturally connects to active learning and OED. Active learning selects informative samples from a given data pool to improve label efficiency (Deng et al., 2023; Settles, 2009; Yang et al., 2023); however, it typically operates reactively over a fixed unlabeled set during training. As a result, the support of the underlying distribution remains fixed as well. We instead study the a priori optimization of the sampling distribution itself. Like our work, there are active learning algorithms that are adaptive to labels (Zhu et al., 2019). However, this approach is pool-based, involves joint optimization instead of bilevel optimization, applies only to classification instead of regression, and adopts a worst-case instead of average-case perspective. On the other hand, OED explicitly reasons about where to make measurements, often through A- or D-optimality or Bayesian and sequential criteria (Atkinson et al., 2007; Fedorov, 2010; Foster et al., 2019; Huan et al., 2024). The classical OED problem emphasizes parameter inference and uncertainty reduction. In contrast, we target distributions that minimize the average OOD prediction error for function and operator learning, a design principle not commonly prioritized in OED. Nevertheless, our work shares similarities with modern measure-centric OED (Jin et al., 2024a), in which the design variable is a probability measure over the design space instead of a fixed number of sample points. In the nonlinear case, measure-based OED requires the formulation of a bilevel optimization problem (Jin et al., 2024b). The difference is that our inner loop involves the supervised training of a machine learning model, while for OED it typically involves solving a PDE-constrained inverse problem.

Other related work includes ridge leverage score and randomly pivoted Cholesky sampling for kernel regression (Rudi et al., 2018; Chen et al., 2025b) as well as Bayesian optimization for sequential point-by-point dataset updates (Shahriari et al., 2015). These techniques provide principled but

typically label-independent sampling rules that do not depend on the target map to be learned. Instead, the present paper proposes adaptive, target-dependent, and supervised distribution design methods.

B ADDITIONAL MATHEMATICAL BACKGROUND

This appendix provides more mathematical background for the present work.

We begin with two comments about Gaussian distributions on Hilbert spaces. First, we remark that the 2-Wasserstein metric between Gaussian measures has a closed-form expression.

Example B.1 (2-Wasserstein metric between Gaussians). Let $\mu = \mathcal{N}(m_{\mu}, \mathcal{C}_{\mu})$ and $\nu = \mathcal{N}(m_{\nu}, \mathcal{C}_{\nu})$ be Gaussian measures on a real separable Hilbert space $(H, \langle \cdot, \cdot \rangle, ||\cdot||)$. Then Gelbrich (1990) gives

$$W_2^2(\mu,\nu) = \|m_{\mu} - m_{\nu}\|^2 + \operatorname{tr}\left(\mathcal{C}_{\mu} + \mathcal{C}_{\nu} - 2\left(\mathcal{C}_{\mu}^{1/2}\mathcal{C}_{\nu}\mathcal{C}_{\mu}^{1/2}\right)^{1/2}\right). \tag{B.1}$$

The square root in (B.1) is the unique self-adjoint positive semidefinite (PSD) square root of a self-adjoint PSD linear operator. The trace of an operator $A\colon H\to H$ is given by $\operatorname{tr}(A)=\sum_{j\in\mathbb{N}}\langle e_j,Ae_j\rangle$ for any orthonormal basis $\{e_j\}_{j\in\mathbb{N}}$ of H. If A is PSD, then we can also write $\operatorname{tr}(A)=\|A^{1/2}\|_{\mathrm{HS}}^2$, where $B\mapsto \|B\|_{\mathrm{HS}}^2:=\sum_{j\in\mathbb{N}}\|Be_j\|^2$ is the squared Hilbert–Schmidt norm.

We frequently invoke (B.1) in this paper. When H is infinite-dimensional, the *Karhunen–Loève* expansion (KLE) is often useful in calculations (Stuart, 2010, Thm. 6.19, p. 533).

Lemma B.2 (KLE of Gaussian measure). Let $m \in (H, \langle \cdot, \cdot \rangle, \| \cdot \|)$ and $\mathcal{C} : H \to H$ be self-adjoint, PSD, and trace-class. Let $\{\lambda_j\}_{j \in \mathbb{N}} \subseteq \mathbb{R}_{\geq 0}$ be a non-increasing rearrangement of the eigenvalues of \mathcal{C} and $\{\phi_j\}_{j \in \mathbb{N}} \subseteq H$ be the corresponding orthonormal set of eigenvectors. If $\{\xi_j\}_{j \in \mathbb{N}}$ is an independent and identically distributed sequence with $\xi_1 \sim \mathcal{N}(0,1)$, then

$$\operatorname{Law}\left(m + \sum_{j \in \mathbb{N}} \sqrt{\lambda_j} \xi_j \phi_j\right) = \mathcal{N}(m, \mathcal{C}).$$
(B.2)

Generalizing Lem. B.2, there exist Karhunen–Loève expansions of non-Gaussian measures. In this case, the $\{\xi_j\}_{j\in\mathbb{N}}$ are non-Gaussian, have mean zero and variance one, and are only pairwise uncorrelated instead of independent and identically distributed (i.i.d.); still, the KLE "diagonalizes" the probability measure.

Next, we turn our attention to general probability measures on abstract spaces. This allows us to define probability measures *over* the space of probability measures.

Definition B.3 (spaces of probability measures). Let (X, \mathcal{B}) be a measurable space. The space of probability measures on X is defined by

$$\mathscr{P}(X) := \{ \mu \colon \mathcal{B} \to [0,1] \mid \mu \text{ is a probability measure} \}.$$

It is equipped with the σ -algebra $\mathcal{B}(\mathscr{P}(X))$ generated by the evaluation maps $e_A \colon \mathscr{P}(X) \to [0,1]$ given by $e_A(\mu) := \mu(A)$ for all $A \in \mathcal{B}$. Similarly, the space of probability measures on $\mathscr{P}(X)$ is

$$\mathscr{P}(\mathscr{P}(X)) \coloneqq \{\mathbb{Q} \colon \mathcal{B}(\mathscr{P}(X)) \to [0,1] \mid \mathbb{Q} \text{ is a probability measure} \}.$$

Now suppose that $(\mathscr{P}(X),\mathsf{d})$ is a metric space. For $p\in[1,\infty)$, we say that $\mathbb{Q}\in\mathscr{P}_p(\mathscr{P}(X))$ if $\mathbb{Q}\in\mathscr{P}(\mathscr{P}(X))$ and $\mathbb{E}_{\mu\sim\mathbb{Q}}[\mathsf{d}^p(\mu,\mu_0)]<\infty$ for some (and hence for all) $\mu_0\in\mathscr{P}(X)$.

Def. B.3 can be related to the notion of random measure from Def. 2.2 as follows. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. Recall that a random probability measure on X is a measurable map

$$\Xi \colon (\Omega, \mathcal{F}) \to (\mathscr{P}(X), \mathcal{B}(\mathscr{P}(X))).$$

Then the law (or distribution) of Ξ is the pushforward measure $\text{Law}(\Xi) \in \mathcal{P}(\mathcal{P}(X))$ defined by

$$\operatorname{Law}(\Xi)(E) := \mathbb{P}(\{\omega \in \Omega \mid \Xi(\omega) \in E\}) \quad \text{for all} \quad E \in \mathcal{B}(\mathscr{P}(X)).$$

In this paper, we typically do not distinguish between a random measure and its law. ¹

We now provide three examples of random measures.

¹Both concepts are closely related to *Markov kernels* (Kallenberg, 2017).

 Example B.4 (empirical measure). Given i.i.d. random variables X_1,\ldots,X_N taking values in measurable space (M,\mathcal{B}) with $X_1 \sim \mu$, the *empirical measure* is the random probability measure $\mu^{(N)} := \frac{1}{N} \sum_{n=1}^N \delta_{X_n}$, where $\delta_x \in \mathscr{P}(M)$ is the Dirac measure at $x \in M$. For each measurable set $A \in \mathcal{B}$, the number $\mu^{(N)}(A) \in [0,1]$ counts the fraction of samples that fall in A. The measure $\mu^{(N)}$ is random because the samples $\{X_n\}_{n=1}^N$ are random.

Example B.5 (Dirichlet process). Let α be a finite nonnegative measure on (M, \mathcal{B}) . A *Dirichlet process* $\mu \sim \mathrm{DP}(\alpha)$ is a random probability measure such that for every k and every measurable partition $\{A_1, \ldots, A_k\}$ of M, it holds that

$$(\mu(A_1), \ldots, \mu(A_k)) \sim \mathsf{Dirichlet}(\alpha(A_1), \ldots, \alpha(A_k))$$
.

Example B.6 (random Gaussian measures). In the infinite-dimensional state space setting, a canonical example is that of a random Gaussian measure defined through the KLE (B.2) of its mean and covariance. Let $(H, \langle \cdot, \cdot \rangle, \| \cdot \|)$ be infinite-dimensional, $\{\phi_j\}$ be an orthonormal basis of H, and $\{\lambda_j\} \subset \mathbb{R}_{\geq 0}$ and $\{\sigma_{ij}\} \subset \mathbb{R}_{\geq 0}$ be summable. Define the random measure $\mathcal{N}(m, \mathcal{C})$ by

$$m \stackrel{\text{Law}}{:=} \sum_{j \in \mathbb{N}} \sqrt{\lambda_j} \xi_j \phi_j \quad \text{and} \quad \mathcal{C} \coloneqq \mathcal{R}^* \mathcal{R}, \quad \text{where} \quad \mathcal{R} \stackrel{\text{Law}}{:=} \sum_{i \in \mathbb{N}} \sum_{j \in \mathbb{N}} \sqrt{\sigma_{ij}} \zeta_{ij} \phi_i \otimes \phi_j \,.$$
 (B.3)

The $\{\xi_j\}$ and $\{\zeta_{ij}\}$ are zero mean, unit variance, pairwise-uncorrelated real random variables. The outer product in (B.3) is defined by $(a\otimes b)c\coloneqq \langle b,c\rangle a$. Since $\mathcal R$ is a KLE in the space of Hilbert–Schmidt operators on H, it follows that $\mathcal C=\mathcal R^*\mathcal R$ is a valid self-adjoint PSD trace-class covariance operator with probability one. Since $\{\phi_j\}$ is fixed, the degrees of freedom in $\mathbb Q:=\mathrm{Law}(\mathcal N(m,\mathcal C))$ are the choice of eigenvalues $\{\lambda_j\}$ and $\{\sigma_{ij}\}$ and the distribution of the coefficients $\{\xi_j\}$ and $\{\zeta_{ij}\}$. Typically the coefficients are chosen i.i.d. from some common distribution such as the standard normal distribution $\mathcal N(0,1)$ or the uniform distribution on the interval $[-\sqrt{3},\sqrt{3}]$. The preceding construction of $\mathcal N(m,\mathcal C)$ immediately extends to mixtures of Gaussian measures.

To conclude **SM** B, we connect back to the barycenters described in Rmk. 4.8.

Definition B.7 (Wasserstein barycenter). Fix $p \in [1, \infty)$. Given a probability measure $\mathbb{Q} \in \mathscr{P}_p(\mathscr{P}_p(\mathcal{U}))$ and a set $P \subseteq \mathscr{P}_p(\mathcal{U})$, a (p, P)-Wasserstein barycenter of \mathbb{Q} is any solution of

$$\inf_{\mu \in \mathsf{P}} \mathbb{E}_{\nu \sim \mathbb{Q}} \, \mathsf{W}_p^p(\mu, \nu) \,. \tag{B.4}$$

If $P = \mathscr{P}_p(\mathcal{U})$, then we simply write p-Wasserstein barycenter for brevity.

When $\mathbb Q$ is supported on Gaussian measures, the 2-Wasserstein barycenter of $\mathbb Q$ can be characterized. **Example B.8** (Gaussian barycenter). Let $\mathbb Q \coloneqq \operatorname{Law}(\mathcal N(m,\mathcal C)) \in \mathscr P_2(\mathscr P_2(\mathcal U))$, where $m \in \mathcal U$ and $\mathcal C \colon \mathcal U \to \mathcal U$ are the random mean vector and random covariance operator of a random Gaussian measure. Then the solution $\overline{\mu}$ to $\inf_{\mu \in \mathscr P_2(\mathcal U)} \mathbb E[W_2^2(\mu,\mathcal N(m,\mathcal C))]$ is the Gaussian $\overline{\mu} = \mathcal N(\mathbb E[m],\overline{\mathcal C})$ (Álvarez-Esteban et al., 2016, Thm. 2.4), where $\overline{\mathcal C}$ is a non-zero PSD solution to the nonlinear equation

$$\overline{C} = \mathbb{E}\left[\left(\overline{C}^{1/2}C\overline{C}^{1/2}\right)^{1/2}\right]. \tag{B.5}$$

One can replace the expectation with an empirical sum and employ fixed-point iterations to approximately solve (B.5) for \overline{C} .

C DETAILS FOR SECTION 3: LIPSCHITZ THEORY FOR OUT-OF-DISTRIBUTION ERROR BOUNDS

This appendix provides the proof of Prop. 3.1 and expands on Cor. 3.2 with Ex. C.1. We begin with the proof.

Proof of Prop. 3.1. We begin by proving the first assertion, which is similar to the result and proof of (Benitez et al., 2024, Lem. 7.2, p. 23). Let $\phi(u) := \|\mathcal{G}_1(u) - \mathcal{G}_2(u)\|$. For any u and u', the triangle inequality yields

$$|\phi(u) - \phi(u')| \le ||\mathcal{G}_1(u) - \mathcal{G}_2(u) - (\mathcal{G}_1(u') - \mathcal{G}_2(u'))||$$

$$\le ||\mathcal{G}_2(u) - \mathcal{G}_2(u')|| + ||\mathcal{G}_1(u) - \mathcal{G}_1(u'))||$$

$$< (\operatorname{Lip}(\mathcal{G}_1) + \operatorname{Lip}(\mathcal{G}_2))||u - u'||.$$

Thus, $\operatorname{Lip}(\phi) \leq \operatorname{Lip}(\mathcal{G}_1) + \operatorname{Lip}(\mathcal{G}_2) =: C$. By Kantorovich–Rubenstein duality, $\mathbb{E}_{u' \sim \nu'} \phi(u') - \mathbb{E}_{u \sim \nu} \phi(u) \leq C W_1(\nu, \nu')$ as asserted.

For the W₂ result, note that $\|\mathcal{G}_i(u)\| \leq \text{Lip}(\mathcal{G}_i)\|u\| + \|\mathcal{G}_i(0)\|$. This and the previous estimate gives

$$\phi^{2}(u') - \phi^{2}(u) = (\phi(u') + \phi(u))(\phi(u') - \phi(u))$$

$$\leq (\|\mathcal{G}_{1}(u)\| + \|\mathcal{G}_{1}(u')\| + \|\mathcal{G}_{2}(u)\| + \|\mathcal{G}_{2}(u')\|)C\|u - u'\|$$

$$\leq (C\|u\| + C\|u'\| + 2\|\mathcal{G}_{1}(0)\| + 2\|\mathcal{G}_{2}(0)\|)C\|u - u'\|.$$

Let π be the W₂-optimal coupling of μ and μ' . By the Cauchy–Schwarz inequality, the expectation $\mathbb{E}_{u'\sim\mu'}\,\phi^2(u')-\mathbb{E}_{u\sim\mu}\,\phi^2(u)=\mathbb{E}_{(u,u')\sim\pi}[\phi^2(u')-\phi^2(u)]$ is bounded above by

$$C\sqrt{\mathbb{E}_{(u,u')\sim\pi} \left(C\|u\| + C\|u'\| + 2\|\mathcal{G}_1(0)\| + 2\|\mathcal{G}_2(0)\|\right)^2} \sqrt{\mathbb{E}_{(u,u')\sim\pi} \|u - u'\|^2}$$

$$\leq C\sqrt{4C^2 \left(\mathsf{m}_2(\mu) + \mathsf{m}_2(\mu')\right) + 16\|\mathcal{G}_1(0)\|^2 + 16\|\mathcal{G}_2(0)\|^2} \, \mathsf{W}_2(\mu,\mu') \,.$$

The last line is due to $(a+b+c+d)^2 \le 4(a^2+b^2+c^2+d^2)$. The second assertion is proved. \Box

Next, we apply Cor. 3.2 in the context of Gaussian mixtures.

Example C.1 (Gaussian mixtures). Consider the *random* Gaussian measure $\nu' = \mathcal{N}(m,\mathcal{C}) \sim \mathbb{Q} \in \mathscr{P}_2(\mathscr{P}_2(\mathcal{U}))$, where $(m,\mathcal{C}) \sim \Pi$ and Π is the mixing distribution over the mean (with marginal Π_1) and covariance (with marginal Π_2). Define the Gaussian mixture $\overline{\nu} := \mathbb{E}_{(m,\mathcal{C}) \sim \Pi} \mathcal{N}(m,\mathcal{C}) \in \mathscr{P}_2(\mathcal{U})$ via duality. Let $\overline{\mu} := \mathbb{E}_{(m',\mathcal{C}') \sim \Pi'} \mathcal{N}(m',\mathcal{C}')$ be another Gaussian mixture with mixing distribution Π' (with marginals denoted by Π'_1 and Π'_2). By linearity of expectation and Fubini's theorem, $\mathscr{E}_{\mathbb{Q}}(\mathcal{G}) = \mathscr{E}_{\delta_{\overline{\nu}}}(\mathcal{G})$; that is, average-case accuracy with respect to \mathbb{Q} is equivalent to accuracy with respect to the single test distribution $\overline{\nu}$; also see (D.2). This fact and Cor. 3.2 yield $\mathscr{E}_{\mathbb{Q}}(\mathcal{G}) \leq \mathbb{E}_{u \sim \overline{\mu}} \|\mathcal{G}^*(u) - \mathcal{G}(u)\|^2 + c(\mathcal{G}^*, \mathcal{G}, \overline{\mu}, \overline{\nu}) \mathsf{W}_2(\overline{\mu}, \overline{\nu})$. The factor $c(\mathcal{G}^*, \mathcal{G}, \overline{\mu}, \overline{\nu})$ can be further simplified using the closed-form expression for second moments of Gaussian measures. For the W₂ factor, Lem. C.2 and coupling arguments imply that

$$\mathscr{E}_{\mathbb{Q}}(\mathcal{G}) \leq \mathbb{E}_{u \sim \overline{\mu}} \left\| \mathcal{G}^{\star}(u) - \mathcal{G}(u) \right\|_{\mathcal{V}}^{2} + c \left(\mathcal{G}^{\star}, \mathcal{G}, \overline{\mu}, \overline{\nu} \right) \sqrt{\mathsf{W}_{2}^{2}(\Pi_{1}, \Pi_{1}^{\prime}) + \mathsf{W}_{2}^{2}(\Pi_{2}, \Pi_{2}^{\prime})} \,. \tag{C.1}$$

The two W₂ distances between the marginals may be further bounded with (B.1) in **SM** B under additional structural assumptions on the mixing distributions; see Rmk. C.3. We often invoke this example in numerical studies with $\mathbb{Q} = \sum_{k=1}^K w_k \delta_{\nu'_k}$ (where $K \in \mathbb{N}$, i.e., a finite mixture), weights $\sum_{k=1}^K w_k = 1$, Gaussian test measures $\{\nu'_k\}$, and degenerate training mixture $\overline{\mu} := \mathcal{N}(m_0, \mathcal{C}_0)$.

The following lemma is used in Ex. C.1; it is an independently interesting result.

Lemma C.2 (Lipschitz stability of mixtures in Wasserstein metric). Let $p \in [1, \infty)$. If \mathbb{Q} and \mathbb{Q}' belong to $\mathscr{P}_p(\mathscr{P}_p(\mathcal{U}))$, then $\mathsf{W}_p(\mathbb{E}_{\mu \sim \mathbb{Q}} \mu, \mathbb{E}_{\nu \sim \mathbb{Q}'} \nu) \leq \mathsf{W}_p(\mathbb{Q}, \mathbb{Q}')$.

Proof. Let $\Pi \in \Gamma(\mathbb{Q}, \mathbb{Q}') \subseteq \mathscr{P}(\mathscr{P}_p(\mathcal{U}) \times \mathscr{P}_p(\mathcal{U}))$ be a W_p-optimal coupling with respect to the ground space $(\mathscr{P}_p(\mathcal{U}), \mathsf{W}_p)$. By the joint convexity of the functional W^p_p and Jensen's inequality,

$$\mathsf{W}_p^p(\mathbb{E}_{\mu \sim \mathbb{Q}} \, \mu, \mathbb{E}_{\nu \sim \mathbb{Q}'} \, \nu) = \mathsf{W}_p^p(\mathbb{E}_{(\mu,\nu) \sim \Pi} \, \mu, \mathbb{E}_{(\mu,\nu) \sim \Pi} \, \nu) \leq \mathbb{E}_{(\mu,\nu) \sim \Pi} \, \mathsf{W}_p^p(\mu,\nu) = \mathsf{W}_p^p(\mathbb{Q},\mathbb{Q}') \, .$$

Taking p-th roots completes the proof.

In Lem. C.2, the Wasserstein distance $W_p(\mathbb{Q}, \mathbb{Q}')$ uses $(\mathscr{P}_p(\mathcal{U}), W_p)$ as the ground metric space in Def. 2.1. Thus, it is a Wasserstein metric over the Wasserstein space.

Last, we remark on structural assumptions that allow us to further bound (C.1) in Ex. C.1.

Remark C.3 (KLE mixing distributions). Instate the setting and notation of Ex. C.1. Furthermore, let $m \sim \Pi_1$ and $m' \sim \Pi'_1$ have centered (possibly non-Gaussian) KLEs

$$m \stackrel{\text{Law}}{=} \sum_{j \in \mathbb{N}} \sqrt{\lambda_j} \xi_j \phi_j \quad \text{and} \quad m' \stackrel{\text{Law}}{=} \sum_{j \in \mathbb{N}} \sqrt{\lambda_j'} \xi_j' \phi_j \,, \quad \text{where} \quad \xi_j \stackrel{\text{i.i.d.}}{\sim} \eta \quad \text{and} \quad \xi_j' \stackrel{\text{i.i.d.}}{\sim} \eta' \,.$$

 In the preceding display, we further use the notation from (B.3) in Ex. B.6; thus, $\{\xi_j\}$ and $\{\xi_j'\}$ have zero mean and unit variance, with laws in $\mathscr{P}_2(\mathbb{R})$. Let $\rho \in \Gamma(\eta, \eta')$ be the W₂-optimal coupling. Similar to (Garbuno-Inigo et al., 2023, Lem. 4.1, pp. 13–14), let $\pi \in \Gamma(\Pi_1, \Pi_1')$ be a coupling with the property that $\pi = \text{Law}(h, h')$, where $h = \sum_{j \in \mathbb{N}} \sqrt{\lambda_j} z_j \phi_j$ and $h' = \sum_{j \in \mathbb{N}} \sqrt{\lambda_j'} z_j' \phi_j$ in law and $(z_j, z_j') \stackrel{\text{i.i.d.}}{\sim} \rho$. Then by (B.1) and the triangle inequality,

$$\begin{aligned} W_{2}^{2}(\Pi_{1}, \Pi_{1}') &\leq \mathbb{E}_{(m, m') \sim \pi} \|m - m'\|^{2} \\ &\leq 2 \,\mathbb{E}_{\rho} \sum_{j \in \mathbb{N}} \lambda_{j} |\xi_{j} - \xi_{j}'|^{2} + 2 \,\mathbb{E}_{\eta'} \sum_{j \in \mathbb{N}} |\xi_{j}'|^{2} \Big| \sqrt{\lambda_{j}} - \sqrt{\lambda_{j}'} \Big|^{2} \\ &= 2 W_{2}^{2}(\eta, \eta') \sum_{j \in \mathbb{N}} \lambda_{j} + 2 \sum_{j \in \mathbb{N}} \Big| \sqrt{\lambda_{j}} - \sqrt{\lambda_{j}'} \Big|^{2}. \end{aligned}$$

This estimate shows that Π_1 and Π_1' are close in W_2 if η and η' are close in W_2 and $\{\sqrt{\lambda_j}\}$ and $\{\sqrt{\lambda_j'}\}$ are close in $\ell^2(\mathbb{N};\mathbb{R})$. We conclude this remark by upper bounding the term $W_2^2(\Pi_2,\Pi_2')$ appearing (C.1). Following (B.3) from Ex. B.6, suppose that $\mathcal{C}=\mathcal{R}^*\mathcal{R}$ and $\mathcal{C}'=(\mathcal{R}')^*\mathcal{R}'$, where

$$\mathcal{R} \stackrel{\text{Law}}{=} \sum_{i \in \mathbb{N}} \sum_{j \in \mathbb{N}} \sqrt{\sigma_{ij}} \zeta_{ij} \phi_i \otimes \phi_j, \, \mathcal{R}' \stackrel{\text{Law}}{=} \sum_{i \in \mathbb{N}} \sum_{j \in \mathbb{N}} \sqrt{\sigma'_{ij}} \zeta'_{ij} \phi_i \otimes \phi_j, \, \zeta_{ij} \stackrel{\text{i.i.d.}}{\sim} \varrho, \, \text{ and } \, \zeta'_{ij} \stackrel{\text{i.i.d.}}{\sim} \varrho'$$

are KLE expansions in the Hilbert space $HS(\mathcal{U})$ of Hilbert–Schmidt operators from \mathcal{U} into itself. Using (B.1) and (Cockayne & Duncan, 2021, Lem. SM1.2, p. SM2), it holds that

$$\begin{aligned} \mathsf{W}_2^2(\Pi_2,\Pi_2') &= \inf_{\gamma \in \Gamma(\Pi_2,\Pi_2')} \mathbb{E}_{(\mathcal{C},\mathcal{C}') \sim \gamma} \, \mathsf{W}_2^2 \big(\mathcal{N}(0,\mathcal{C}), \mathcal{N}(0,\mathcal{C}') \big) \\ &\leq \inf_{\gamma \in \Gamma(\Pi_2,\Pi_2')} \mathbb{E}_{\gamma} \big\| \mathcal{R} - \mathcal{R}' \big\|_{\mathrm{HS}}^2 \,. \end{aligned}$$

Then a nearly identical coupling argument to the one used for $W_2(\Pi_1, \Pi'_1)$ delivers the bound

$$\mathsf{W}_2^2(\Pi_2,\Pi_2') \leq 2\mathsf{W}_2^2(\varrho,\varrho') \sum_{i\in\mathbb{N}} \sum_{j\in\mathbb{N}} \sigma_{ij} + 2\sum_{i\in\mathbb{N}} \sum_{j\in\mathbb{N}} \left| \sqrt{\sigma_{ij}} - \sqrt{\sigma_{ij}'} \right|^2.$$

D DETAILS FOR SECTION 4: PRACTICAL ALGORITHMS FOR TRAINING DISTRIBUTION DESIGN

This appendix expands on the content of Sec. 4 and provides the deferred proofs of the theoretical results from that section. It begins in **SM** D.1 by identifying a degeneracy of the bilevel problem formulation when infinite data are available. **SM** D.2 concerns the bilevel minimization algorithm from Sec. 4.1 and **SM** D.3 pertains to the alternating minimization algorithm from Sec. 4.2.

D.1 NECESSITY OF FINITE DATA

Although Secs. 3 and 4 are formulated at the infinite data level, i.e., with full access to $\mathbb Q$ and the candidate training distribution ν , the proposed formulation is only advantageous when ν is accessible from a finite number of samples. Of course, this is the case in practice (and moreover $\mathbb Q$ will often only be accessible from samples as well); recall Rem. 4.1. Our *optimize-then-discretize* approach of Secs. 3 and 4 simplifies the exposition and notation, while allowing us to derive continuum algorithms that still work at the finite data level. Regardless, it is instructive to see how the necessity of finite data emerges from a theoretical viewpoint.

To this end, recall the functional $\mathcal{E}_{\mathbb{Q}}$ from (3.3). Define the infinite mixture probability measure

$$\nu_{\mathbb{Q}} := \mathbb{E}_{\nu' \sim \mathbb{Q}}[\nu'] \tag{D.1}$$

via duality with bounded continuous test functions. By linearity, one can check that

$$\mathcal{E}_{\mathbb{Q}}(\mathcal{G}) = \mathbb{E}_{u' \sim \nu_{\mathbb{Q}}} \| \mathcal{G}^{\star}(u') - \mathcal{G}(u') \|_{\mathcal{Y}}^{2}. \tag{D.2}$$

This display shows that the average OOD accuracy with respect to \mathbb{Q} is equal to the standard test error with respect to the test mixture distribution $\nu_{\mathbb{Q}}$.

Now define the optimal value $V_{\mathbb{Q}}^{\star} \coloneqq \inf_{\mathcal{G} \in \mathcal{H}} \mathcal{E}_{\mathbb{Q}}(\mathcal{G})$. By the definition of the trained model in (4.1), it holds that $V_{\mathbb{Q}}^{\star} = \mathcal{E}_{\mathbb{Q}}(\widehat{\mathcal{G}}^{(\nu_{\mathbb{Q}})})$. The inequality $V_{\mathbb{Q}}^{\star} \le \inf_{\nu} \mathcal{E}_{\mathbb{Q}}(\widehat{\mathcal{G}}^{(\nu)})$ holds because $V_{\mathbb{Q}}^{\star}$ is the optimal value over all elements from \mathcal{H} . Moreover, $\inf_{\nu} \mathcal{E}_{\mathbb{Q}}(\widehat{\mathcal{G}}^{(\nu)}) \le \mathcal{E}_{\mathbb{Q}}(\widehat{\mathcal{G}}^{(\nu_{\mathbb{Q}})}) = V_{\mathbb{Q}}^{\star}$ due to optimality over training distributions. We deduce that

$$\inf_{\nu} \mathcal{E}_{\mathbb{Q}}(\widehat{\mathcal{G}}^{(\nu)}) = \inf_{\mathcal{G} \in \mathcal{H}} \mathcal{E}_{\mathbb{Q}}(\mathcal{G}). \tag{D.3}$$

That is, the optimal value of the objective over distributions versus the objective over models coincide. Write $\widehat{\nu}$ for a solution of $\inf_{\nu} \mathcal{E}_{\mathbb{Q}}(\widehat{\mathcal{G}}^{(\nu)})$ from (4.1). The preceding display then says

$$\mathcal{E}_{\mathbb{Q}}(\widehat{\mathcal{G}}^{(\widehat{\nu})}) = \mathcal{E}_{\mathbb{Q}}(\widehat{\mathcal{G}}^{(\nu_{\mathbb{Q}})}). \tag{D.4}$$

This does not imply that $\widehat{\nu}$ equals $\nu_{\mathbb{Q}}$. Indeed, these infinite-dimensional optimization problems over possibly noncompact sets could admit multiple minima. What (D.4) does suggest is that the mixture $\nu_{\mathbb{Q}}$ (which only depends on \mathbb{Q}) is an equally good solution for making the OOD accuracy small compared to the adaptive measure $\widehat{\nu}$ (which depends on \mathcal{G}^{\star} , \mathcal{H} , and \mathbb{Q}), assuming full knowledge of \mathbb{Q} and infinite training data. In this case, there is no reason to construct $\widehat{\nu}$ if $\nu_{\mathbb{Q}}$ is known.

However, for each N, writing $\nu^N = \frac{1}{N} \sum_{n=1}^N \delta_{u_n}$ for the empirical measure corresponding to i.i.d. samples $u_n \sim \nu$, we can instead consider

$$\widehat{\nu}_{N} \in \underset{\nu \in \mathbb{P} \subset \mathscr{P}_{2}(\mathcal{U})}{\arg \min} \left\{ \mathscr{E}_{\mathbb{Q}}(\widehat{\mathcal{G}}_{N}^{(\nu)}) \, \middle| \, \widehat{\mathcal{G}}_{N}^{(\nu)} \in \underset{\mathcal{G} \in \mathcal{H}}{\arg \min} \, \mathbb{E}_{u \sim \nu^{N}} \big\| \mathcal{G}^{\star}(u) - \mathcal{G}(u) \big\|_{\mathcal{Y}}^{2} \right\}$$
(D.5)

for some search space P. The preceding argument leading to the degenerate behavior shown in (D.4) no longer holds for the empirical bilevel problem (D.5) due to the finite data constraint. It is now possible that models trained on samples from $\widehat{\nu}_N$ can outperform models trained on samples from the mixture $\nu_{\mathbb{Q}}$. Indeed, we observe the superiority of $\widehat{\nu}_N$ over $\nu_{\mathbb{Q}}$ even for simple Gaussian parametrizations in Secs. 5 and F.1. More generally, the *test distribution* $\nu_{\mathbb{Q}}$ is not the optimal distribution from which to generate training data for regression problems. This fact is known to the numerical analysis community (Adcock, 2024; Cohen & Migliorati, 2017), building off of variance reduction and importance sampling ideas (Agapiou et al., 2017) for Monte Carlo and quasi-Monte Carlo methods (Caflisch, 1998; Dick et al., 2013). It remains true for (linear) operator learning, where de Hoop et al. (2023) suggests that one should pick a training data distribution over input functions whose samples are rougher than function samples from the test distribution. This implies that the training distribution support should (strictly) contain the support of the test distribution. However, there are fundamental limits on how much improvement one can expect to obtain by adjusting the training data distribution (Grohs et al., 2025; Kovachki et al., 2024a).

Another practical takeaway from the preceding discussions is that if data from $\mathbb Q$ are very limited and fixed once and for all, then a finite sample approximation will not capture the mixture $\nu_{\mathbb Q}$ well. This could lead to a poor performing model when trained on the aggregate empirical mixture samples. On the other hand, the proposed approach delivers a distribution $\widehat{\nu}_N$ that ideally can be sampled from at will (leading to many more training samples, e.g., a generative model) and uses information about \mathcal{G}^* and \mathcal{H} to compensate for limited information about $\mathbb Q$.

D.2 Derivations for subsection 4.1: Exact bilevel formulation

Proofs. Recall the use of the adjoint method in real (i.e., $\mathcal{Y} = \mathbb{R}$) RKHS \mathcal{H} with kernel κ from Sec. 4.1. For a fixed $\nu \in \mathscr{P}_2(\mathcal{U})$, the data misfit $\Psi \colon \mathcal{H} \to \mathbb{R}$ for this ν is

$$\mathcal{G} \mapsto \Psi(\mathcal{G}) := \frac{1}{2} \| \iota_{\nu} \mathcal{G} - \mathcal{G}^{\star} \|_{L_{\nu}^{2}(\mathcal{U})}^{2} = \mathbb{E}_{u \sim \nu} |\mathcal{G}(u) - \mathcal{G}^{\star}(u)|^{2}.$$
 (D.6)

This is the inner objective in (4.1). In the absence of RKHS norm penalization, typically one must assume the existence of minimizers of Ψ (Rudi et al., 2015, Assump. 1, p. 3). The next lemma characterizes them.

Lemma D.1 (inner minimization solution). Fix a set $P \subseteq \mathscr{P}_2(\mathcal{U})$ of probability measures. Suppose that for every $\nu \in P$, there exists $\widehat{\mathcal{G}}^{(\nu)} \in \mathcal{H}$ such that $\Psi(\widehat{\mathcal{G}}^{(\nu)}) = \inf_{\mathcal{G} \in \mathcal{H}} \Psi(\mathcal{G})$, where Ψ is as in (D.6). Then $\widehat{\mathcal{G}}^{(\nu)}$ satisfies the compact operator equation $\mathcal{K}_{\nu}\widehat{\mathcal{G}}^{(\nu)} = \iota_{\nu}^*\mathcal{G}^*$ in \mathcal{H} .

Proof. The result follows from convex optimization in separable Hilbert space. Since we assume minimizers of Ψ exist, they are critical points. The Frechét derivative $D\Psi: \mathcal{H} \to \mathcal{H}^*$ is

$$\mathcal{G} \mapsto D\Psi(\mathcal{G}) = \langle \iota_{\nu}^*(\iota_{\nu}\mathcal{G} - \mathcal{G}^*), \cdot \rangle_{\kappa}.$$

At a critical point \mathcal{G} , we have $D\Psi(\mathcal{G})=0$ in \mathcal{H}^* . This means

$$\langle \iota_{\nu}^*(\iota_{\nu}\mathcal{G} - \mathcal{G}^*), h \rangle_{\kappa} = 0 \quad \text{for all} \quad h \in \mathcal{H}.$$

By choosing $h = \iota_{\nu}^*(\iota_{\nu}\mathcal{G} - \mathcal{G}^*)$, it follows that $\iota_{\nu}^*\iota_{\nu}\mathcal{G} = \mathcal{K}_{\nu}\mathcal{G} = \iota_{\nu}^*\mathcal{G}^*$ in \mathcal{H} as asserted.

Next, recall the Lagrangian \mathcal{L} from (4.2). We prove Lem. 4.2 on the adjoint state equation.

Proof of Lem. 4.2. The action of the Frechét derivative with respect to the second coordinate is

$$D_2 \mathcal{L}(\nu, \mathcal{G}, \lambda) = \mathbb{E}_{\nu' \sim \mathbb{Q}} \langle \iota_{\nu'}^* (\iota_{\nu'} \mathcal{G} - \mathcal{G}^*), \cdot \rangle_{\kappa} + \langle \mathcal{K}_{\nu} \lambda, \cdot \rangle_{\kappa} \in \mathcal{H}^*.$$

Setting the preceding display equal to $0 \in \mathcal{H}^*$, we deduce that

$$\mathcal{K}_{\nu}\lambda = \mathbb{E}_{\nu' \sim \mathbb{Q}} \left[\iota_{\nu'}^* (\mathcal{G}^* - \iota_{\nu'} \mathcal{G}) \right] \quad \text{in} \quad \mathcal{H}.$$

Simplifying the right-hand side using the definition of ι_{ν}^* completes the proof.

We are now in a position to prove Prop. 4.3, which computes the first variation of $\mathcal{J}: \nu \mapsto \mathcal{L}(\nu, \widehat{\mathcal{G}}^{(\nu)}, \lambda^{(\nu)}(\widehat{\mathcal{G}}^{(\nu)}))$ in the sense of (Santambrogio, 2015, Def. 7.12, p. 262).

Proof of Prop. 4.3. The first variation (Santambrogio, 2015, Def. 7.12, p. 262) of \mathcal{L} with respect to the variable ν is given by the function

$$D_1 \mathcal{L}(\nu, \mathcal{G}, \lambda) = (\mathcal{G} - \mathcal{G}^*)\lambda$$

that maps from \mathcal{U} to \mathbb{R} , which is independent of ν given \mathcal{G} and λ . Now let $\widehat{\mathcal{G}}^{(\nu)}$ be as in Lem. D.1 and $\lambda^{(\nu)}$ be as in Lem. 4.2. Due to the smoothness of \mathcal{L} (4.2) with respect to both \mathcal{G} and ν —in particular, \mathcal{L} is convex in \mathcal{G} and linear in ν —the envelope theorem (Afriat, 1971) ensures that

$$(D\mathcal{J})(\nu) = (D_1 \mathcal{L})\big|_{(\nu, \mathcal{G}, \lambda) = (\nu, \widehat{\mathcal{G}}^{(\nu)}, \lambda^{(\nu)}(\widehat{\mathcal{G}}^{(\nu)}))} = (\widehat{\mathcal{G}}^{(\nu)} - \mathcal{G}^{\star})\lambda^{(\nu)}(\widehat{\mathcal{G}}^{(\nu)})$$

as asserted.

Remark D.2 (Wasserstein gradient). As discussed in Rmk. 4.4, it is clear that

$$\nabla \big[(D\mathcal{J})(\nu) \big] = \big(\nabla \widehat{\mathcal{G}}^{(\nu)} - \nabla \mathcal{G}^{\star} \big) \lambda^{(\nu)} \big(\widehat{\mathcal{G}}^{(\nu)} \big) + \big(\widehat{\mathcal{G}}^{(\nu)} - \mathcal{G}^{\star} \big) \nabla \big(\lambda^{(\nu)} \big(\widehat{\mathcal{G}}^{(\nu)} \big) \big)$$

is the Wasserstein gradient of $\mathcal J$ (Santambrogio, 2015, Chp. 8.2). Its computation requires the usual $\mathcal U$ -gradient of the trained model $\widehat{\mathcal G}^{(\nu)}$, the true map $\mathcal G^\star$, and the optimal adjoint state $\lambda^{(\nu)}(\widehat{\mathcal G}^{(\nu)})$.

We now turn to the parametric training distribution setting. Let $\nu = \nu_{\vartheta}$, where $\vartheta \in \mathcal{V}$ and $(\mathcal{V}, \langle \cdot, \cdot \rangle_{\mathcal{V}}, \| \cdot \|_{\mathcal{V}})$ is a parameter Hilbert space as introduced in Sec. 4.1. The main result of Sec. 4.1 is Thm. 4.5. Before proving that theorem, we require the following lemma.

Lemma D.3 (differentiation under the integral). Let the density $p_{\vartheta} := d\nu_{\vartheta}/d\mu_0$ be such that $(u, \vartheta) \mapsto p_{\vartheta}(u)$ is sufficiently regular. Then for all bounded continuous functions $\phi \in C_b(\mathcal{U})$ and all $\vartheta \in \mathcal{V}$, it holds that

$$\nabla_{\vartheta} \int_{\mathcal{U}} \phi(u) \nu_{\vartheta}(du) = \int_{\mathcal{U}} \phi(u) \Big(\nabla_{\vartheta} \log p_{\vartheta}(u) \Big) \nu_{\vartheta}(du) \,. \tag{D.7}$$

Proof. By the assumed smoothness and absolute continuity with respect to μ_0 ,

$$\nabla_{\vartheta} \int_{\mathcal{U}} \phi(u) \nu_{\vartheta}(du) = \int_{\mathcal{U}} \phi(u) \nabla_{\vartheta} p_{\vartheta}(u) \mu_0(du).$$

Shrinking the domain of integration to the support of ν_{ϑ} , multiplying and dividing the rightmost equality by $p_{\vartheta}(u)$, and using the logarithmic derivative formula delivers the assertion.

With the preceding lemma in hand, we are now able to prove Thm. 4.5.

Proof of Thm. 4.5. We use the chain rule for Frechét derivatives to compute $DJ: \mathcal{V} \to \mathcal{V}^*$ and hence the gradient $\nabla J: \mathcal{V} \to \mathcal{V}$. Define $Z: \mathcal{V} \to \mathscr{P}_2(\mathcal{U})$ by $Z(\vartheta) = \nu_{\vartheta}$. Then Lem. D.3 and duality allow us to identify DZ with the (possibly signed) \mathcal{V} -valued measure

$$(DZ(\vartheta))(du) = (\nabla_{\vartheta} \log p_{\vartheta}(u))Z(\vartheta)(du).$$

We next apply the chain rule for Frechét derivatives to the composition $J = \mathcal{J} \circ Z$. This yields

$$DJ(\vartheta) = (D\mathcal{J})(Z(\vartheta)) \circ (DZ)(\vartheta) \in \mathcal{V}^*$$

for any $\vartheta \in \mathcal{V}$. To interpret this expression, let $\vartheta' \in \mathcal{V}$ be arbitrary. Then formally,

$$\begin{split} \left(D\mathsf{J}(\vartheta)\right)(\vartheta') &= \int_{\mathcal{U}} \left(\widehat{\mathcal{G}}^{(Z(\vartheta))}(u) - \mathcal{G}^{\star}(u)\right) \left(\lambda^{(Z(\vartheta))}(\widehat{\mathcal{G}}^{(Z(\vartheta))})\right)(u) \left\langle \left(DZ(\vartheta)\right)(du), \vartheta'\right\rangle_{\mathcal{V}} \\ &= \left\langle \int_{\mathcal{U}} \left(\widehat{\mathcal{G}}^{(\nu_{\vartheta})}(u) - \mathcal{G}^{\star}(u)\right) \left(\lambda^{(\nu_{\vartheta})}(\widehat{\mathcal{G}}^{(\nu_{\vartheta})})\right)(u) \left(\nabla_{\vartheta} \log p_{\vartheta}(u)\right) \nu_{\vartheta}(du), \vartheta'\right\rangle_{\mathcal{V}}. \end{split}$$

The formula (4.4) for ∇J follows.

Discretized implementation. A practical, discretized implementation of the parametric bilevel gradient descent scheme Alg. 1 further requires regularization of compact operator equations and finite data approximations. When the inner objective of (4.1) is approximated over a finite data set, we have the following regularized solution for the model update obtained from kernel ridge regression. The lemma uses the compact vector and matrix notation $U = (u_1, \dots, u_N)^{\top} \in \mathcal{U}^N$, $\mathcal{G}^{\star}(U)_n = \mathcal{G}^{\star}(u_n)$ for every n, and $\kappa(U, U)_{ij} = \kappa(u_i, u_j)$ for every i and j.

Lemma D.4 (model update). Fix a regularization strength $\sigma > 0$. Let $\{u_n\}_{n=1}^N \sim \nu^{\otimes N}$. Then

$$\widehat{\mathcal{G}}_{N}^{(\nu,\sigma)} := \sum_{n=1}^{N} \kappa(\cdot, u_{n}) \beta_{n} \approx \widehat{\mathcal{G}}^{(\nu)}, \quad \text{where} \quad \beta = \left(\kappa(U, U) + N\sigma^{2} I_{N}\right)^{-1} \mathcal{G}^{\star}(U) \in \mathbb{R}^{N}. \quad (D.8)$$

Proof. To numerically update \mathcal{G} by solving $\mathcal{K}_{\nu}\mathcal{G} = \iota_{\nu}^*\mathcal{G}^*$ in \mathcal{H} for fixed ν , we first regularize the inverse with a small nugget parameter $\sigma > 0$. Then we define

$$\widehat{\mathcal{G}}^{(\nu,\sigma)} := (\mathcal{K}_{\nu} + \sigma^2 \operatorname{Id}_{\mathcal{H}})^{-1} \iota_{\nu}^* \mathcal{G}^* = \iota_{\nu}^* (\iota_{\nu} \iota_{\nu}^* + \sigma^2 \operatorname{Id}_{L_{\nu}^2})^{-1} \mathcal{G}^* \approx \widehat{\mathcal{G}}^{(\nu)}.$$

The second equality is due to the usual Woodbury push-through identity for compact linear operators. Let $\nu^{(N)} \coloneqq \frac{1}{N} \sum_{n=1}^N \delta_{u_n}$ be the empirical measure with $u_n \overset{\text{i.i.d.}}{\sim} \nu$. We identify $L^2_{\nu^{(N)}}$ with \mathbb{R}^N equipped with the usual (i.e., unscaled) Euclidean norm. The action of $\mathcal{K}_{\nu^{(N)}}$ on a function $h \in \mathcal{H}$ is

$$\mathcal{K}_{\nu^{(N)}}h = \frac{1}{N} \sum_{n=1}^{N} \kappa(\cdot, u_n) h(u_n) =: \frac{1}{N} \kappa(\cdot, U) h(U)$$
 (D.9)

and similar for $\iota_{\nu^{(N)}}^*$ acting on \mathbb{R}^N . Define $\widehat{\mathcal{G}}_N^{(\nu,\sigma)} \coloneqq \widehat{\mathcal{G}}^{(\nu^{(N)},\sigma)}$. By the preceding displays,

$$\widehat{\mathcal{G}}_N^{(\nu,\sigma)} = \sum_{n=1}^N \kappa(\,\cdot\,,u_n)\beta_n\,, \quad \text{where} \quad \beta = \left(\kappa(U,U) + N\sigma^2 I_N\right)^{-1}\mathcal{G}^\star(U)\,,$$

which recovers standard kernel ridge regression under square loss and Gaussian likelihood. This implies the asserted approximation $\widehat{\mathcal{G}}_N^{(\nu,\sigma)} \approx \widehat{\mathcal{G}}^{(\nu)}$.

In Lem. D.4, we expect that $\widehat{\mathcal{G}}_N^{(\nu,\sigma)} \to \widehat{\mathcal{G}}^{(\nu)}$ as $N \to \infty$ and $\sigma \to 0$. This lemma takes care of the finite data approximation of the model. Next, we tackle the adjoints.

When discretizing the integral in the gradient (4.4) with Monte Carlo sampling over the training data, we recognize that only the values of $\lambda = \lambda^{(\nu)}(\widehat{\mathcal{G}}^{(\nu)})$ at the training points are needed and not $\lambda \in \mathcal{H}$ itself. That is, we never need to query λ away from the current training data samples. The next lemma exploits this fact to numerically implement the adjoint state update.

Lemma D.5 (adjoint state update). For $J \in \mathbb{N}$ and a sequence $\{M_j\} \subseteq \mathbb{N}$, define $\mathbb{Q}^{(J)} := \frac{1}{J} \sum_{j=1}^{J} \delta_{\mu_j^{(M_j)}}$, where $\mu_j^{(M_j)} := \frac{1}{M_j} \sum_{m=1}^{M_j} \delta_{v_m^j}$, $V^j := \{v_m^j\}_{m=1}^{M_j} \sim \mu_j^{\otimes M_j}$, and $\{\mu_j\}_{j=1}^{J} \sim \mathbb{Q}^{\otimes J}$. Then with $\widehat{\mathcal{G}}_N^{(\nu,\sigma)}$ and U as in (D.8), the adjoint values $(\lambda^{(\nu)}(\widehat{\mathcal{G}}^{(\nu)}))(U) \in \mathbb{R}^N$ are approximated by

$$\widehat{\lambda}_{N,J}^{(\nu,\sigma)}(U) \coloneqq \Big(\kappa(U,U) + N\sigma^2 I_N\Big)^{-1} \frac{N}{J} \sum_{j=1}^J \frac{1}{M_j} \kappa(U,V^j) \Big(\mathcal{G}^\star(V^j) - \widehat{\mathcal{G}}_N^{(\nu,\sigma)}(V^j)\Big) \,. \tag{D.10}$$

Proof. Let $\mu \sim \mathbb{Q}$, where $\mathbb{Q} \in \mathscr{P}_2(\mathscr{P}_2(\mathcal{U}))$ is fixed. For some M, let $\mu^{(M)} = \frac{1}{M} \sum_{m=1}^M \delta_{v_m}$ be the empirical measure with $V = \{v_m\}_{m=1}^M \sim \mu^{\otimes M}$. Consider the case J=1 so that $\mathbb{Q}^{(1)} = \delta_{\mu^{(M)}}$; the general case for J>1 as asserted in the lemma follows by linearity. With $\nu=\nu^{(N)}$ being the empirical data measure corresponding to $U\in\mathcal{U}^N$, $\mathcal{G}=\widehat{\mathcal{G}}_N^{(\nu,\sigma)}$ being the trained model from Lem. D.4, and $\mathbb{Q}^{(1)}$ in place of \mathbb{Q} , the adjoint equation (4.3) from Lem. 4.2 is

$$\mathcal{K}_{\nu^{(N)}}\lambda = \frac{1}{M} \sum_{m=1}^{M} \kappa(\cdot, v_m) \left(\mathcal{G}^{\star}(v_m) - \widehat{\mathcal{G}}_N^{(\nu, \sigma)}(v_m) \right).$$

By evaluating the preceding display at the N data points $U = \{u_n\}_{n=1}^N$ and using (D.9), we obtain

$$\frac{1}{N}\kappa(U,U)\lambda(U) = \frac{1}{M}\kappa(U,V) \left(\mathcal{G}^{\star}(V) - \widehat{\mathcal{G}}_{N}^{(\nu,\sigma)}(V)\right).$$

This is a linear equation for the unknown $\lambda(U) \in \mathbb{R}^N$. Finally, we regularize the equation by replacing $\kappa(U,U)$ with $\kappa(U,U) + \sigma^2 I_N$. We denote the solution of the resulting system by $\widehat{\lambda}_{N,1}^{(\nu,\sigma)}(U)$, which is our desired approximation (D.10) of $(\lambda^{(\nu)}(\widehat{\mathcal{G}}^{(\nu)}))(U) \in \mathbb{R}^N$.

In Lem. D.5, the vector $\widehat{\lambda}_{N,J}^{(\nu,\sigma)}(U) \in \mathbb{R}^N$ depends on ν (through the empirical measure $\nu^{(N)} \coloneqq \frac{1}{N} \sum_{n=1}^N \delta_{u_n}$), the regularization parameter σ , the current model $\widehat{\mathcal{G}}_N^{(\nu,\sigma)}$, and the random empirical measure $\mathbb{Q}^{(J)}$ (which itself depends on the sequence $\{M_j\}_{j=1}^J$). The formula (D.10) requires the inversion of the same regularized kernel matrix as in the model update (D.8). Thus, it is advantageous to pre-compute a factorization of this matrix so that it may be reused in the adjoint update step.

Finally, we perform the numerically implementable training data distribution update by inserting the approximations $\widehat{\mathcal{G}}_N^{(\nu,\sigma)}$ and $\widehat{\lambda}_{N,J}^{(\nu,\sigma)}(U)$ from Lems. D.4–D.5 and empirical measure $\nu_{\vartheta}^{(N)}$ into the exact gradient formula (4.4) for $\nabla \mathsf{J}$. This yields the approximate gradient

$$\widehat{\nabla} \mathsf{J}(\vartheta) \coloneqq \int_{\mathcal{U}} (\widehat{\mathcal{G}}_{N}^{(\nu_{\vartheta},\sigma)}(u) - \mathcal{G}^{\star}(u)) \widehat{\lambda}_{N,J}^{(\nu_{\vartheta},\sigma)}(u) (\nabla_{\vartheta} \log p_{\vartheta}(u)) \nu_{\vartheta}^{(N)}(du), \qquad (D.11)$$

where $\nu_{\vartheta}^{(N)}$ is the empirical measure of the same training data points used to obtain both the trained model and adjoint state vector, and we identify the vector $\widehat{\lambda}_{N,J}^{(\nu,\sigma)}(U) \in \mathbb{R}^N$ with a function $\widehat{\lambda}_{N,J}^{(\nu,\sigma)} \in L^2_{\nu_{\vartheta}^{(N)}}$. This leads to Alg. 3. In this practical algorithm, we allow for a varying number N_k of training samples from $\nu_{\vartheta^{(k)}}$ and a varying regularization strength σ_k at each iteration k in Line 4. The integrals in (D.11) and Line 4 of Alg. 3 are with respect to the empirical measure; hence, they equal the equally-weighted average over the finite data samples that make up the empirical measure.

One limitation of the practical bilevel gradient descent scheme Alg. 3 as written is that it is negatively impacted by overfitting in the model update step, as discussed in the following remark.

Remark D.6 (overfitting and vanishing gradients). A drawback of approximating the integral in the exact gradient (4.4) with a Monte Carlo average over the i.i.d. training samples in (D.11) and Line 4 of Alg. 3 is that the resulting gradient approximation incorporates the training error residual. If the trained model overfits to the data, then this residual becomes extremely small, which slows down the convergence of gradient descent significantly. In the extreme case, the gradient vanishes if the model perfectly fits the training data. Thus, the proposed method as currently formulated would likely fail for highly overparametrized models. To address this problem, one solution is to hold out a "validation" subset of the training data pairs for the adjoint state and Monte Carlo integral calculations;

Algorithm 3 Gradient Descent on a Parametrized Bilevel Objective: Discretized Scheme

- 1: **Initialize:** Parameter $\vartheta^{(0)}$, step sizes $\{t_k\}$, nuggets $\{\sigma_k\}$, sample sizes $\{N_k\}$, measure $\mathbb{Q}^{(J)}$
- 2: **for** $k = 0, 1, 2, \dots$ **do**

- 3: **Sample and label:** Construct the empirical measure $\nu_{\vartheta^{(k)}}^{(N_k)}$ and label vector $\mathcal{G}^{\star}(U) \in \mathbb{R}^{N_k}$
- 4: **Gradient step:** Update the training distribution's parameters via

$$\vartheta^{(k+1)} = \vartheta^{(k)} - t_k \int_{\mathcal{U}} \Big(\widehat{\mathcal{G}}_{N_k}^{(\nu_{\vartheta^{(k)}}, \sigma_k)}(u) - \mathcal{G}^{\star}(u) \Big) \widehat{\lambda}_{N_k, J}^{(\nu_{\vartheta^{(k)}}, \sigma_k)}(u) \Big(\nabla_{\vartheta} \log p_{\vartheta^{(k)}}(u) \Big) \nu_{\vartheta^{(k)}}^{(N_k)}(du) \Big)$$

- 5: **if** stopping criterion is met **then**
- 6: Return $\vartheta^{(k+1)}$ then **break**

the remaining data should be used to train the model. Another approach is to normalize the size of the approximate gradient by its current magnitude. This adaptive normalization aims to avoid small gradients. In the numerical experiments of Sec. 5, we use kernel ridge regression with strictly positive regularization; this prevents interpolation of the data and helps avoid the vanishing gradient issue.

D.3 DERIVATIONS FOR SUBSECTION 4.2: UPPER-BOUND MINIMIZATION

Recall the alternating minimization algorithm Alg. 2 from Sec. 4.2. We provide details about its numerical implementation. For gradient-based implementations, the first variation of the upper bound objective (4.8) is required. This is the content of Prop. 4.9, which we now prove.

Proof of Prop. 4.9. Since the first term $\int_{\mathcal{U}} f(u)\nu(du)$ in the objective F in (4.8) is linear in the argument ν , the first variation of that term is simply f itself (Santambrogio, 2015, Def. 7.12, p. 262). The chain rule of Euclidean calculus and a similar first variation calculation delivers the second term in the asserted derivative (4.9). The third and final term arising from the chain rule requires the first variation of $\nu \mapsto \mathbb{E}_{\nu' \sim \mathbb{Q}} W_2^2(\nu, \nu')$. By (Santambrogio, 2015, Prop. 7.17, p. 264, applied with cost $c(x,y) = \|x-y\|_{\mathcal{U}}^2/2$), this first variation equals the Kantorovich potential $2\phi_{\nu,\nu'}$, where the asserted characterization of $\phi_{\nu,\nu'}$ is due to Brenier's theorem; see (Chewi, 2024, Thm. 1.3.8, p. 27) or (Villani, 2021). Finally, we exchange expectation over \mathbb{Q} and differentiation using dominated convergence due to the assumed regularity of the set \mathbb{P} .

The Wasserstein gradient of F is useful for interacting particle discretizations. We have the following corollary of Prop. 4.9.

Corollary D.7 (Wasserstein gradient). *Instate the hypotheses of Prop. 4.9. The Wasserstein gradient* ∇DF at $\nu \in P \subseteq \mathscr{P}_2(\mathcal{U})$ is given by the function mapping any $u \in \mathcal{U}$ to the vector

$$\big((\nabla DF)(\nu) \big)(u) = \nabla f(u) + \sqrt{\frac{\mathbb{E}_{\nu' \sim \mathbb{Q}} \operatorname{W}_2^2(\nu, \nu')}{1 + \operatorname{m}_2(\nu)}} u + \sqrt{\frac{1 + \operatorname{m}_2(\nu)}{\mathbb{E}_{\nu' \sim \mathbb{Q}} \operatorname{W}_2^2(\nu, \nu')}} \big(u - \mathbb{E}_{\nu' \sim \mathbb{Q}} T_{\nu \to \nu'}(u) \big).$$

Alternatively, one can derive standard gradients when working with a parametrized family $\{\nu_{\vartheta}\}_{\vartheta\in\mathcal{V}}$ of candidate training distributions. The next lemma instantiates this idea in the specific setting of a finite-dimensional Gaussian family.

Lemma D.8 (chain rule: Gaussians). *Instate the hypotheses of Prop.* 4.9. Suppose that $\mathcal{U} \subseteq \mathbb{R}^d$. Let $\nu_{\vartheta} \coloneqq \mathcal{N}(m_{\vartheta}, \mathcal{C}_{\vartheta})$ and $\mathsf{F}(\vartheta) \coloneqq F(\nu_{\vartheta})$, where F is given in (4.8). Write $p_{\vartheta} \colon \mathcal{U} \to \mathbb{R}_{\geq 0}$ for the density of ν_{ϑ} with respect to any σ -finite dominating measure. Then for any parameter $\vartheta \in \mathcal{V}$, it

holds that

$$\nabla \mathsf{F}(\vartheta) = \int_{\mathcal{U}} f(u) \left(\nabla_{\vartheta} \log p_{\vartheta}(u) \right) \nu_{\vartheta}(du)$$

$$+ \frac{1}{2} \sqrt{\frac{\mathbb{E}_{\nu' \sim \mathbb{Q}} \mathsf{W}_{2}^{2}(\nu_{\vartheta}, \nu')}{1 + \|m_{\vartheta}\|_{\mathcal{U}}^{2} + \operatorname{tr}(\mathcal{C}_{\vartheta})}} \int_{\mathcal{U}} \frac{\|u\|_{\mathcal{U}}^{2}}{2} \left(\nabla_{\vartheta} \log p_{\vartheta}(u) \right) \nu_{\vartheta}(du)$$

$$+ \sqrt{\frac{1 + \|m_{\vartheta}\|_{\mathcal{U}}^{2} + \operatorname{tr}(\mathcal{C}_{\vartheta})}{\mathbb{E}_{\nu' \sim \mathbb{Q}} \mathsf{W}_{2}^{2}(\nu_{\vartheta}, \nu')}} \int_{\mathcal{U}} \frac{1}{2} \left\langle u, (I_{d} - \mathbb{E}_{\mathbb{Q}}[A'_{\vartheta}]) u \right\rangle_{\mathcal{U}} \left(\nabla_{\vartheta} \log p_{\vartheta}(u) \right) \nu_{\vartheta}(du)$$

$$+ \sqrt{\frac{1 + \|m_{\vartheta}\|_{\mathcal{U}}^{2} + \operatorname{tr}(\mathcal{C}_{\vartheta})}{\mathbb{E}_{\nu' \sim \mathbb{Q}} \mathsf{W}_{2}^{2}(\nu_{\vartheta}, \nu')}} \int_{\mathcal{U}} \left\langle u, \mathbb{E}_{\mathbb{Q}}[A'_{\vartheta}] m_{\vartheta} - \mathbb{E}_{\mathbb{Q}}[m'] \right\rangle_{\mathcal{U}} \left(\nabla_{\vartheta} \log p_{\vartheta}(u) \right) \nu_{\vartheta}(du) ,$$

where $\nu' = \mathcal{N}(m', \mathcal{C}') \sim \mathbb{Q}$ is a random Gaussian and $A'_{\vartheta} := \mathcal{C}_{\vartheta}^{-1/2} (\mathcal{C}_{\vartheta}^{1/2} \mathcal{C}' \mathcal{C}_{\vartheta}^{1/2})^{1/2} \mathcal{C}_{\vartheta}^{-1/2}$.

Proof. The W₂-optimal transport map from $\nu = \mathcal{N}(m_1, \mathcal{C}_1)$ to $\nu' = \mathcal{N}(m_2, \mathcal{C}_2)$ is

$$u \mapsto T_{\nu \to \nu'}(u) := m_2 + A(u - m_1), \text{ where } A := \mathcal{C}_1^{-1/2} \left(\mathcal{C}_1^{1/2} \mathcal{C}_2 \mathcal{C}_1^{1/2}\right)^{1/2} \mathcal{C}_1^{-1/2}$$

Furthermore, $\phi_{\nu,\nu'} = \|\cdot\|_{\mathcal{U}}^2/2 - \varphi_{\nu,\nu'}$, where $\nabla \varphi_{\nu,\nu'} = T_{\nu \to \nu'}$ by Brenier's theorem (Chewi, 2024, Thm. 1.3.8, p. 27). Upon integration, we find that

$$u \mapsto \phi_{\nu,\nu'}(u) = \frac{\|u\|^2}{2} - \frac{1}{2}\langle u, Au \rangle - \langle u, m_2 - Am_1 \rangle.$$

This formula is unique up to a constant. Now let $\nu := \nu_{\vartheta} = \mathcal{N}(m_{\vartheta}, \mathcal{C}_{\vartheta})$ and $\nu' = \mathcal{N}(m', \mathcal{C}') \sim \mathbb{Q}$ be as in the hypotheses of the lemma. By the chain rule for Frechét derivatives,

$$\nabla \mathsf{F}(\vartheta) = \int_{\mathcal{U}} (DF(\nu_{\vartheta}))(u) (\nabla_{\vartheta} \log p_{\vartheta}(u)) \nu_{\vartheta}(du)$$

for each ϑ . Applying (4.9) from Prop. 4.9, closed-form formulas for the second moments of Gaussian distributions, and the Fubini–Tonelli theorem to exchange integrals completes the proof.

Lem. D.8 requires the gradient $\nabla_{\vartheta} \log p_{\vartheta}$ of the log density with respect to the parameter ϑ . In general, this gradient has no closed form and must be computed numerically, e.g., with automatic differentiation. However, when only the mean of the Gaussian family is allowed to vary, the following remark discusses an analytical expression for the gradient.

Remark D.9 (Gaussian mean parametrization). Let $\mathcal{U} \subseteq \mathbb{R}^d$. If $\vartheta = m$ so that $\nu_{\vartheta} = \nu_m = \mathcal{N}(m, \mathcal{C}_0)$ for fixed \mathcal{C}_0 , then $\nabla_m \log p_m(u) = \mathcal{C}_0^{-1}(u-m)$, where p_m is the Lebesgue density of ν_m .

E OPERATOR LEARNING ARCHITECTURES

In this appendix, we provide detailed descriptions of the various deep operator learning architectures used in the paper. These are the DeepONet (SM E.1), NIO (SM E.2), and the newly proposed AMINO (SM E.3).

E.1 DEEPONET: DEEP OPERATOR NETWORK

Deep operator networks (DeepONet) (Lu et al., 2021) provide a general framework for learning nonlinear operators \mathcal{G}^* : $u \mapsto \mathcal{G}^*(u)$ from data. A DeepONet \mathcal{G}_{θ} achieves this via two jointly trained subnetworks:

- Branch network B: takes as input the values of the input function u at a fixed set of sensor locations $\{x^{(j)}\}_{j=1}^{J}$ and outputs a feature vector $b \in \mathbb{R}^p$ for some fixed latent dimension p.
- Trunk network T: takes as input a query point z in the output domain and outputs feature vector t(z) ∈ ℝ^p.

The operator prediction at function u and query point z is then formed by the inner product of these features, yielding

$$\mathcal{G}_{\theta}(u)(z) = \left\langle \mathsf{B}(u(x^{(1)}), \dots, u(x^{(J)})), \mathsf{T}(z) \right\rangle_{\mathbb{R}^p} = \sum_{k=1}^p \mathsf{B}_k(u(x^{(1)}), \dots, u(x^{(J)})) \mathsf{T}_k(z). \quad (E.1)$$

In practice, a constant bias vector is added to the preceding display. During ERM training as in (2.1), the weights θ of both neural networks B and T are optimized together to minimize the average discrepancy between $\mathcal{G}_{\theta}(u)(z)$ and ground truth values $\mathcal{G}^{\star}(u)(z)$ over a dataset of input-output function pairs. This branch-trunk architecture enables DeepONet to approximate a broad class of operators, including the Neumann-to-Dirichlet (NtD) map in EIT and the conductivity-to-solution map in Darcy flow considered in this work; see Sec. 5 and SM F.

E.2 NIO: NEURAL INVERSE OPERATOR

Traditional approaches for constructing direct solvers for inverse problems, such as the D-bar method for EIT (Siltanen et al., 2000), have laid the foundation for direct solver approximations utilizing modern machine learning techniques. Among these, the Neural Inverse Operator (NIO) (Molinaro et al., 2023) framework stands out as a significant advancement. In the context of operator learning, many inverse problems can be formulated as a mapping from an operator Λ_a , dependent on a physical parameter a, to the parameter a itself, i.e., $\Lambda_a \mapsto a$. In practical scenarios, the operator Λ_a is not directly accessible; instead, we observe its action on a set of input-output function pairs $\{(f_n,g_n)\}_n$, where $g_n = \Lambda_a(f_n)$. This observation leads to a two-step problem: first, infer the operator Λ_a from the data pairs, and second, deduce the parameter a from the inferred operator. This process can be further generalized by considering the operator as a pushforward map $(\Lambda_a)_\#$, which transforms an input distribution μ of functions to an output distribution ν_a of functions, i.e., $(\Lambda_a)_\#\mu = \nu_a$; see Sec. 2 for the definition of the pushforward operation. Consequently, a measure-centric machine learning approach to solving the inverse problem can be reformulated as learning the mapping $(\mu, \nu_a) \mapsto a$ (Nelsen & Yang, 2025, Sec. 2.4.1).

Conditioned on a fixed input distribution μ , NIO is designed to learn mappings from output measures ν_a to functions a, i.e., $\nu_a \mapsto a$, effectively relaxing the operator-to-function problem to a distribution-to-function problem. Since direct manipulation of measures is computationally infeasible, NIO circumvents this by leveraging finite data approximations. Specifically, it utilizes a composition of DeepONets and Fourier Neural Operators (FNO) (Li et al., 2021) to approximate the inverse mapping from an empirical measure/approximation $\widehat{\nu}_a$ to the underlying parameter function a.

E.3 AMINO: A MEASURE-THEORETIC INVERSE NEURAL OPERATOR

Despite the innovative architecture of NIO, it exhibits certain limitations, particularly concerning its handling of input distributions. NIO primarily focuses on the output distribution ν_a during evaluation, neglecting variability and influence of the input distribution μ . This oversight can lead to models that perform well on in-distribution (ID) data but lack robustness when faced with OOD scenarios. To address this shortcoming, we introduce *A Measure-theoretic Inverse Neural Operator* (AMINO), a framework that explicitly accounts for the input distribution in the inverse problem. By formulating the inverse task in a fully measure-theoretic setting and taking into account the input distribution, AMINO aims to improve OOD accuracy. Like NIO, AMINO combines a DeepONet and an FNO. The key distinction is that AMINO incorporates both input and output functions as concatenated pairs $\{(f_n,g_n)\}_n \overset{\text{i.i.d.}}{\sim} (\text{Id},\Lambda_a)_\# \mu$ which are passed to the branch network, as illustrated in Fig. 6. See (Nelsen & Yang, 2025, Sec. 4.2.3) for a more detailed description of NIO and AMINO.

NIO versus AMINO. To compare the ID and OOD error between NIO and AMINO, we start with the following forward problem setup taken from Molinaro et al. (2023) for EIT on the domain $\Omega := (0,1)^2 \subset \mathbb{R}^2$. The model is the two-dimensional elliptic PDE boundary value problem

$$\begin{cases} \nabla \cdot (a\nabla u) = 0 & \text{in } \Omega, \\ u = f & \text{on } \partial\Omega. \end{cases}$$
 (E.2)

Figure 6: Schematic of the AMINO architecture, a variant of NIO, which combines a DeepONet, averaging, and an FNO module. The model takes as input the paired function samples $\{(f_\ell,g_\ell)\}_{\ell=1}^L$ drawn from the joint distribution $(\mathrm{Id},\Lambda_a)_\#\mu$ and outputs the target parameter a.

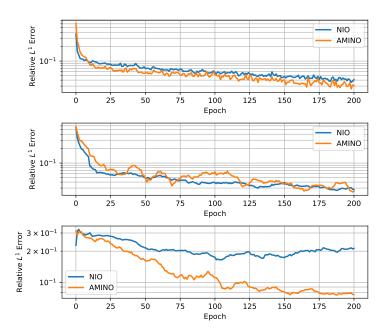


Figure 7: Training and testing performance of AMINO and NIO. (Top): Training error vs. iterations. (Middle) In-distribution testing error with $\omega \sim \text{Unif}([5\pi,\frac{11\pi}{2}])$, matching the training distribution performance. (Bottom) Out-of-distribution testing error with $\omega \sim \text{Unif}([\frac{\pi}{2},7\pi])$, differing from the training distribution. A rolling average with a window of 10 epochs is applied to the middle and bottom plots. AMINO demonstrates improved out-of-distribution performance compared to that of the NIO framework.

For one conductivity realization a defined by

$$x \mapsto a(x) = \exp\left(\sum_{k=1}^{m} c_k \sin(k\pi x_1) \sin(k\pi x_2)\right),\tag{E.3}$$

where $m = \lfloor \overline{m} \rfloor$, $\overline{m} \sim \mathsf{Unif}([1,5])$, and $\{c_k\} \sim \mathsf{Unif}([-1,1]^m)$, we take $\mathsf{L} \in \mathbb{N}$ different Dirichlet boundary conditions $\{f_\ell\}_{\ell=1}^\mathsf{L}$ for f in (E.2). These are given by

$$x \mapsto f_{\ell}(x) = \cos\left(\omega\left(x_1\cos(\theta_{\ell}) + x_2\sin(\theta_{\ell})\right)\right)$$
 (E.4)

with $\omega \sim \text{Unif}([\pi,3\pi])$ and $\theta_l \coloneqq 2\pi\ell/\text{L}$ for $\ell=1,\ldots,\text{L}$. We denote by Λ_a the DtN operator that maps the Dirichlet data of the PDE, f_ℓ , to the Neumann boundary data, $a\frac{\partial u}{\partial \mathbf{n}}|_{\partial\Omega}$, where u is the PDE solution and \mathbf{n} is the outward normal unit vector. Here, the Dirichlet boundary conditions are functions drawn from a probability distribution μ . AMINO learns the map from the joint Dirichlet–Neumann distribution $(\mathrm{Id},\Lambda_a)_\#\mu$ to the conductivity a. In contrast, NIO only takes into account the associated Neumann data distribution $(\Lambda_a)_\#\mu$ as its input, with the entire neural operator conditioned on a fixed Dirichlet data distribution μ .

To compare ID generalization, both models are trained and tested on data sampled from the same Dirichlet distributions. The first two panels in Fig. 7 show that AMINO and NIO perform comparably in both training and ID settings. For OOD evaluation, the models are trained on Dirichlet data

with $\omega \sim \text{Unif}([5\pi, \frac{11\pi}{2}])$ and tested on data with $\omega \sim \text{Unif}([\frac{\pi}{2}, 7\pi])$. The bottom panel in Fig. 7 demonstrates that AMINO achieves significantly better OOD accuracy than NIO, reducing the OOD testing error by approximately half in this example.

F DETAILS FOR SECTION 5: NUMERICAL RESULTS

This appendix describes detailed experimental setups for all numerical tests from Sec. 1 and Sec. 5 in the main text. Regarding function approximation, **SM** F.1 covers the bilevel gradient descent experiments. The remaining appendices concern operator learning. **SM** F.2 details the AMINO experiments. **SM** F.3 and **SM** F.4 describe the results of the alternating minimization algorithm when applied to learning linear NtD maps and the nonlinear solution operator of Darcy flow, respectively. Last, **SM** F.5 presents results from applying nonparametric particle-based AMA to learn a parameter-to-solution map for the radiative transport equation.

F.1 BILEVEL MINIMIZATION FOR FUNCTION APPROXIMATION

We now describe the setting of the numerical results reported in Sec. 5 for Alg. 1 and its numerical version Alg. 3. The goal of the experiments is to find the optimal Gaussian training distribution $\nu_{\vartheta} = \mathcal{N}(m, \mathcal{C}) \in \mathscr{P}_2(\mathbb{R}^d)$ with respect to the bilevel OOD objective (4.1), specialized to the task of function approximation. We optimize over the mean and covariance $\vartheta = (m, \mathcal{C})$. For all $x \in \mathbb{R}^d$, the target test functions $\mathcal{G}^* = g_i : \mathbb{R}^d \to \mathbb{R}$ for $i \in \{1, 2, 3, 4\}$ that we consider are given by

$$g_1(x) = \prod_{j=1}^d \frac{|4x_j - 2| + (j-2)/2}{1 + (j-2)/2},$$
 (F.1a)

$$g_2(x) = 10\sin(\pi x_1 x_2) + 20(x_3 - 1/2)^2 + 10x_4 + 5x_5,$$
 (F.1b)

$$g_3(x) = \sqrt{(100x_1)^2 + \left(x_3(520\pi x_2 + 40\pi) - \frac{1}{(520\pi x_2 + 40\pi)(10x_4 + 1)}\right)^2}, \text{ and } (F.1c)$$

$$g_4(x) = \sum_{l=1}^{1000} c_l \kappa(x, x_l) \,. \tag{F.1d}$$

The first function comes from (Dunbar et al., 2025) and the middle two from (Potts & Schmischke, 2021; Saha et al., 2023). They are common benchmarks in the approximation theory literature; g_1 is called the Sobol G function, g_2 the Friedmann-1 function, and g_3 the Friedmann-2 function. Although these functions are typically defined on compact domains (e.g., $[0,1]^d$ or $[-\pi,\pi]^d$), we instead consider approximating them on average with respect to Gaussian distributions (or mixtures thereof) on the whole of \mathbb{R}^d . The final function g_4 is a linear combination of kernel sections for some kernel function $\kappa\colon \mathbb{R}^d\times\mathbb{R}^d\to\mathbb{R}$. We sample (and then fix) the coefficients $c_l\overset{\text{i.i.d.}}{\sim}$ Unif([-1,1]) and the centers $x_l\overset{\text{i.i.d.}}{\sim}$ Unif([-4,4]) $^{\otimes d}$. We allow the dimension d to be arbitrary for all test functions; in particular, the functions g_2 and and g_3 are constant in certain directions for large enough d.

Model architecture and optimization algorithm. In all bilevel minimization experiments, we use the squared exponential kernel function $\kappa\colon\mathbb{R}^d\times\mathbb{R}^d\to\mathbb{R}_{>0}$ defined by

$$(x, x') \mapsto \kappa(x, x') := \exp\left(-\frac{\|x - x'\|_2^2}{\ell^2}\right),$$
 (F.2)

where $\ell>0$ is a scalar lengthscale hyperparameter. While not tuned for accuracy, we do adjust ℓ on the order of O(1) to O(10) depending on the target function g_i to avoid blow up of the bilevel gradient descent iterations. In particular, we set $\ell=1$ for g_1 , $\ell=3$ for g_2 , $\ell=2/1.1$ for g_3 , and $\ell=5$ for g_4 . We also experimented with the less regular Laplace kernel $(x,x')\mapsto \exp(-\|x-x'\|_1/\ell)$; although giving similar accuracy as κ , the Laplace kernel led to slower bilevel optimization. Thus, we opted to report all results using the squared exponential kernel κ . This same kernel is used to define g_4 in (F.1), which belongs to the RKHS of κ as a result. Given a candidate set of training samples, the resulting kernel method is trained on these samples via kernel ridge regression as in

Lem. D.4. We solve all linear systems with lower triangular direct solves using Cholesky factors of the regularized kernel matrices. No iterative optimization algorithms are required to train the kernel regressors. When fitting the kernel models to data outside of the bilevel gradient descent loop (see the center and right columns of Fig. 8, for instance), we set the ridge strength to be $\sigma^2 = 10^{-3}/N$, where N is the current training sample size.

Training and optimization procedure. Moving on to describe the specific implementation details of Alg. 3, we initialize $\vartheta^{(0)} = (m_0, I_d)$. We take $m_0 = \frac{1}{2}(1, \dots, 1)^{\top} \in \mathbb{R}^d$ in all setups except for the one concerning test function g_1 , where we set $m_0 = (0, \dots, 0)^{\top}$. We decrease the gradient descent step sizes $\{t_k\}$ in Alg. 3 according to a cosine annealing scheduler (Loshchilov & Hutter, 2017) with initial learning rate of 10^{-2} and final learning rate of 0. The regularization parameters $\{\sigma_k\}$ are also selected according to the cosine annealing schedule with initial value 10^{-3} and final value 10^{-7} . Although the sequence $\{N_k\}$ of training sample sizes could also be scheduled, we opt to fix $N_k = 250$ for all k in Alg. 3 for simplicity. No major differences were observed when choosing $\{N_k\}$ according to some cosine annealing schedule. The identification of more principled ways to select the sequences $\{\sigma_k\}$, $\{t_k\}$, and $\{N_k\}$ to optimally balance cost and accuracy is an important direction for future work.

For the probability measure $\mathbb{Q} \in \mathscr{P}_2(\mathscr{P}_2(\mathbb{R}^d))$ over test distributions, we choose it to be the empirical measure

$$\mathbb{Q} := \frac{1}{K} \sum_{k=1}^{K} \delta_{\nu'_k} \,. \tag{F.3}$$

The test measures $\nu_k' \in \mathscr{P}_2(\mathbb{R}^d)$ in (F.3) are deterministic and fixed. In applications, these could be the datasets for which we require good OOD accuracy. However, to synthetically generate the $\{\nu_k'\}$ in practice, we set

$$\nu_k' \coloneqq \mathcal{N}(m_k', \mathcal{C}_k') \,, \quad \text{where} \quad m_k' \overset{\text{i.i.d.}}{\sim} \, \mathcal{N}(0, I_d) \quad \text{and} \quad \mathcal{C}_k' \overset{\text{i.i.d.}}{\sim} \, \text{Wishart}_d(I_d, d+1) \qquad \text{(F.4)}$$

for each $k=1,\ldots,K$. Once the K realizations of these distributions are generated, we view them as fixed once and for all. We assume that K and $\mathbb Q$ are fixed, but we only have access to the empirical measure $\widehat{\mathbb Q}$ defined by (F.3) but with each ν_k' replaced by the empirical measure $\frac{1}{M}\sum_{m=1}^M \delta_{v_m^{(k)}}$ generated by samples $\{v_m^{(k)}\}_{m=1}^M \sim \nu_k'^{\otimes M}$. In the experiments, we set M=5000 and either K=10 or K=1 (a single test distribution). To create a fair test set, we use a fixed 500 sample validation subset to update the adjoint state as in Lem. D.5 and to monitor the validation OOD accuracy during gradient descent. The remaining 4500 samples are used to evaluate the final OOD test accuracy of our trained models; see (F.5).

Based on the available samples $\{v_m^{(k)}\}_{1 \leq m \leq M, \ 1 \leq k \leq K}$ associated to the meta test distribution \mathbb{Q} , we construct four nonadaptive distributions for comparison:

- 1. (Normal) The Gaussian measure $\mathcal{N}(m_0, I_d)$ which is independent of \mathbb{Q} ;
- 2. (Barycenter) The empirical W_2 barycenter of \mathbb{Q} (which is Gaussian), as described in Ex. B.8;
- 3. (Mixture) The non-Gaussian empirical mixture of $\mathbb Q$ given by $\frac{1}{K}\sum_{k=1}^K \mathcal N(\widehat m_k,\widehat{\mathcal C}_k)$, where $\widehat m_k$ and $\widehat{\mathcal C}_k$ are the empirical mean and covariance of the points $\{v_m^{(k)}\}_{m=1}^M$, respectively. Note that we enforce the mixture components to be Gaussian instead of empirical measures;
- 4. (Uniform) The d-dimensional uniform distribution $\mathrm{Unif}([0,1]^d) \coloneqq \mathrm{Unif}([0,1])^{\otimes d}$, which is independent of \mathbb{Q} .

It remains to describe the optimization of the Gaussian family $\nu_{\vartheta} = \mathcal{N}(m,\mathcal{C})$ with parameter $\vartheta = (m,\mathcal{C})$. The closed form update for the mean follows from Rmk. D.9. For the covariance matrix, we optimize over the lower triangular Cholesky factor L of $\mathcal{C} = LL^{\top} \in \mathbb{R}^{d \times d}$ instead of \mathcal{C} itself. This removes redundant parameters due to the symmetry of \mathcal{C} . We enforce strict positive definiteness of \mathcal{C} by replacing all nonpositive diagonal entries of L with a threshold value of 10^{-7} . Other than this constraint, each entry of L is optimized over the whole of \mathbb{R} . Recalling the log density formula (4.5) from Ex. 4.6, we use automatic differentiation to compute $\nabla_L \log(p_{(m,LL^{\top})})$; this factor appears in the definition of the approximate gradient (D.11) and in Alg. 3.

Model evaluation. We evaluate the performance of all kernel regressors $\widehat{\mathcal{G}}_N^{(\nu,\sigma)}$ regularized with strength σ^2 and trained with N i.i.d. samples from ν with respect to the root relative average OOD squared error, which is defined as

$$\mathsf{Err} \coloneqq \sqrt{\frac{\mathbb{E}_{\widehat{\nu'} \sim \widehat{\mathbb{Q}}} \, \mathbb{E}_{u' \sim \widehat{\nu'}} \big| \mathcal{G}^{\star}(u') - \widehat{\mathcal{G}}_{N}^{(\nu,\sigma)}(u') \big|^{2}}{\mathbb{E}_{\widehat{\nu'} \sim \widehat{\mathbb{Q}}} \, \mathbb{E}_{u' \sim \widehat{\nu'}} \big| \mathcal{G}^{\star}(u') \big|^{2}}} \approx \sqrt{\frac{\mathscr{E}_{\mathbb{Q}}(\widehat{\mathcal{G}}_{N}^{(\nu,\sigma)})}{\mathscr{E}_{\mathbb{Q}}(0)}}} \,. \tag{F.5}$$

In the preceding display, we observe that Err employs a Monte Carlo discretization of the integrals appearing in $\mathscr{E}_{\mathbb{Q}}$ by averaging over the empirical meta measure $\widehat{\mathbb{Q}}$, which is composed of K test measures. Each individual test measure is itself an empirical measure with 4500 atoms at points that are unseen during the adjoint state calculation. See the preceding paragraph for details about our choice of \mathbb{Q} . We report values of Err in Figs. 2, 8, 10, 12, 14 and in Tab. 1.

Supplementary numerical results are given in Figs. 8–15. Fig. 8 reports the same information that Fig. 2 does, except now for the remaining functions $\{g_i\}$. The conclusions from these numerical results are mostly the same. The left column shows the training history over 50 iterations; the "seen" curve computes the OOD error with the $\mathbb Q$ samples used in the bilevel algorithm, while the "unseen" curve computes the OOD error on a held-out set of $\mathbb Q$ samples. The third row shows a slight "generalization gap" for $\mathcal G^\star = g_2$ (d=8). The optimized Gaussian outperforms alternative distributions except for the g_3 (d=5) and g_4 cases, where the empirical $\mathbb Q$ mixture eventually overtakes it. Fig. 10 shows that this behavior is largely mitigated if 1000 iterations of gradient descent are used to find the optimal Gaussian instead of 50. We remark that in Fig. 2, the leftmost subplot showing the training history is designed to highlight the long-time behavior (i.e., 1000 iterations) of the bilevel algorithm. In practice, we only run the algorithm for O(10) to O(100) gradient descent iterations; this is reflected in the leftmost column of Figs. 8 and 12.

While the optimized Gaussian is shown to consistently outperform nonadaptive distributions when \mathbb{Q} in (F.3) is composed of K=10 atoms (Figs. 8 and 10), this is less true for a single test distribution (K=1), as seen in Figs. 12 and 14 when N is sufficiently large. In this case, the barycenter and mixture distributions are the same test distribution and lead to trained models with lower OOD error. When N is small and many iterations are used to find the optimal Gaussian, the optimal Gaussian can achieve smaller OOD error than models trained on the test distribution. Understanding for what \mathcal{G}^{\star} , \mathbb{Q} , N, and gradient descent iteration count makes this desirable improvement possible is an interesting and important question for future work.

To understand what properties the optimal Gaussian absorbs from \mathcal{G}^* and \mathbb{Q} , Figs. 9, 11, 13, and 15 visualize the learned covariance matrices, barycenter covariances, and mixture covariances. For instance, the learned covariance puts high weight on the x_3 coordinate in the g_2 case (Fig. 11, rows 2–3, diagonal entry of matrix) even though the barycenter and mixture do not. This makes sense because g_2 is quadratic in x_3 and lower order in all other variables (recall (F.1)); thus, the learned covariance picks out the highest order term in the ground truth map. Similarly, for g_3 the learned covariance assigns high importance to the x_2 and x_3 coordinates which are also fast growing. In all cases (except for g_4), the learned covariance differs substantially from the mixture or test distribution's covariance. These covariance visualization figures highlight an additional benefit of our \mathcal{G}^* -dependent estimators: they lead to interpretable training distributions tailored to the structure of the problem.

Computational resources. All bilevel minimization experiments are performed in Python using the PyTorch framework in float32 single precision on a machine equipped with an NVIDIA GeForce RTX 4090 GPU (24 GB VRAM) running CUDA version 12.4, Intel i7-10700K CPU (8 cores/16 threads) running at 4.9 GHz, and 64 GB of DDR4-3200 CL16 RAM. The experiments required less than one total CPU and GPU hour to complete. An additional five CPU and GPU hours were expended to tune hyperparameters and set up the experiments. We ensure reproducibility of our numerical results by appropriately setting PyTorch backends to be deterministic and using fixed seeds for the relevant random number generators.

F.2 AMINO FOR LEARNING ELECTRICAL IMPEDANCE TOMOGRAPHY

In all machine learning tasks, the choice of training distribution plays a critical role. In practice, testing distributions can span a wide range, making training on the entire space infeasible. Instead,

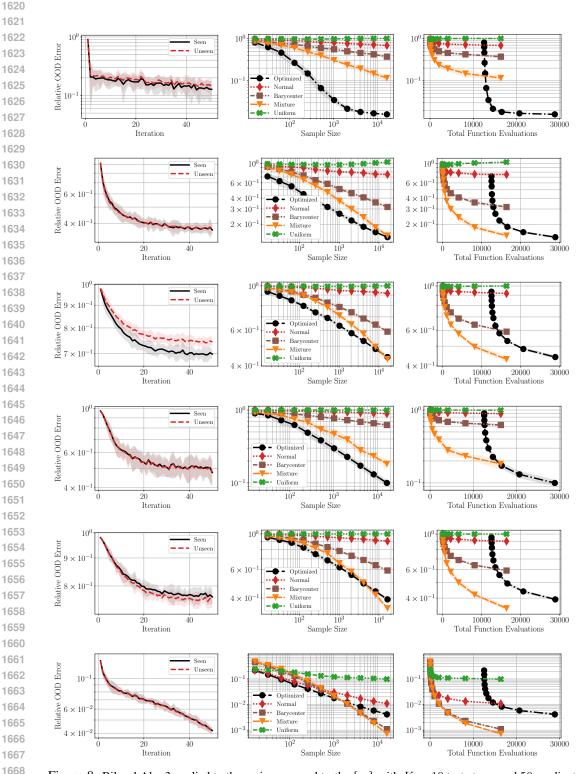


Figure 8: Bilevel Alg. 3 applied to the various ground truths $\{g_i\}$ with K=10 test atoms and 50 gradient descent iterations. (Left column) Evolution of Err from (F.5) vs. iterations. (Center column) Err of model trained on N samples from optimal ν_{ϑ} vs. nonadaptive distributions. (Right column) Same as center, except incorporating the additional function evaluation cost incurred from Alg. 1. From top to bottom, each row represents target functions $\{g_i\}$ corresponding to those listed in each column of Tab. 1 (from left to right). Shading represents two standard deviations away from the mean Err over 10 independent runs.

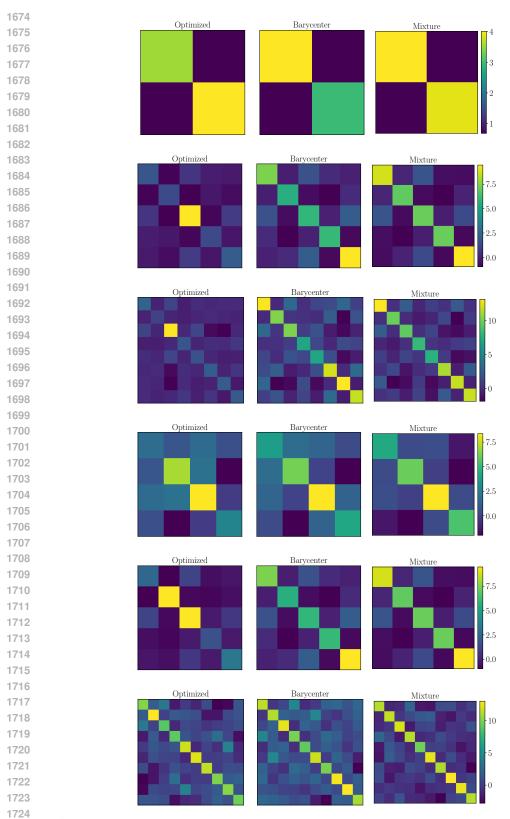


Figure 9: Bilevel Alg. 3 applied with K=10 test atoms and 50 gradient descent iterations. Each column visualizes the covariance matrix of the optimized (left), empirical \mathbb{Q} W₂ barycenter (center), and empirical \mathbb{Q} mixture distributions (right). From top to bottom, each row represents target functions $\{g_i\}$ corresponding to those listed in each column of Tab. 1 (from left to right).

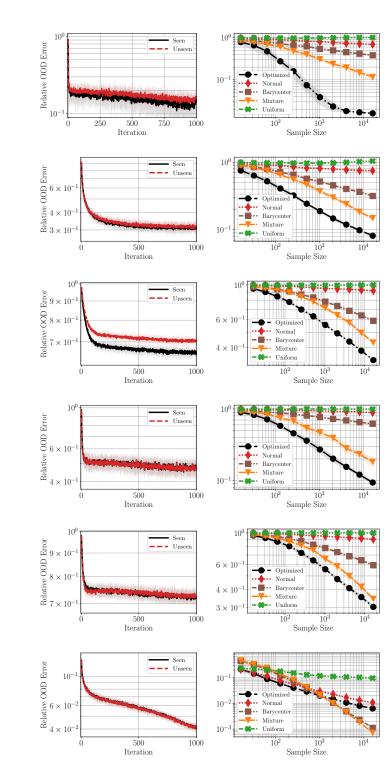


Figure 10: Bilevel Alg. 3 applied to the various ground truths $\{g_i\}$ with K=10 test atoms and 1000 gradient descent iterations. (Left column) Evolution of Err from (F.5) vs. iterations. (Center column) Err of model trained on N samples from optimal ν_{ϑ} vs. nonadaptive distributions. (Right column) Same as center, except incorporating the additional function evaluation cost incurred from Alg. 1. From top to bottom, each row represents target functions $\{g_i\}$ corresponding to those listed in each column of Tab. 1 (from left to right). Shading represents two standard deviations away from the mean Err over 10 independent runs.

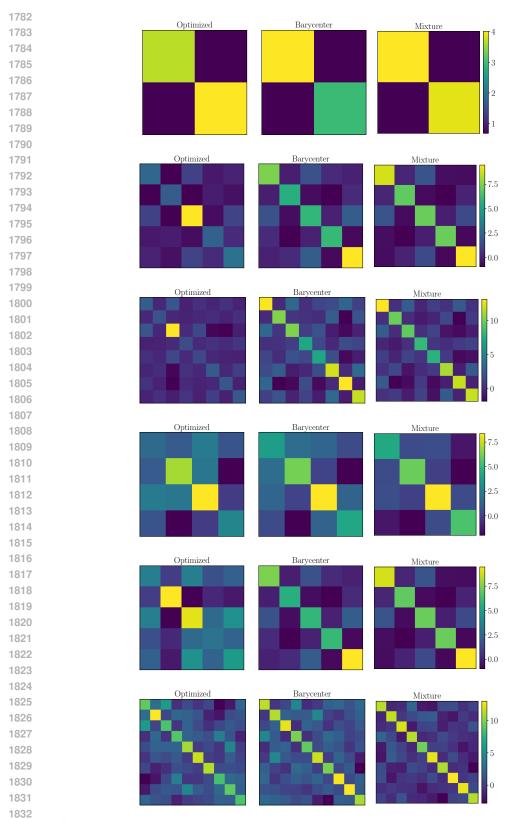


Figure 11: Bilevel Alg. 3 applied with K=10 test atoms and 1000 gradient descent iterations. Each column visualizes the covariance matrix of the optimized (left), empirical \mathbb{Q} W₂ barycenter (center), and empirical \mathbb{Q} mixture distributions (right). From top to bottom, each row represents target functions $\{g_i\}$ corresponding to those listed in each column of Tab. 1 (from left to right).

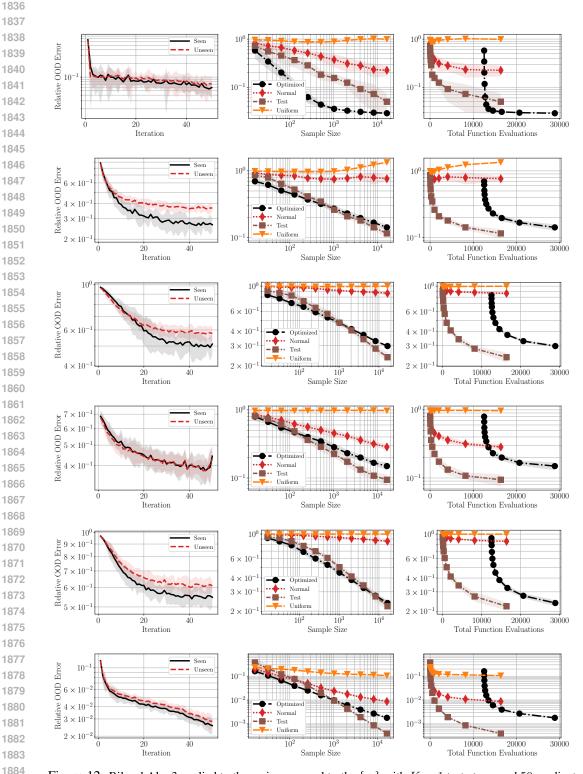


Figure 12: Bilevel Alg. 3 applied to the various ground truths $\{g_i\}$ with K=1 test atoms and 50 gradient descent iterations. (Left column) Evolution of Err from (F.5) vs. iterations. (Center column) Err of model trained on N samples from optimal ν_{ϑ} vs. nonadaptive distributions. (Right column) Same as center, except incorporating the additional function evaluation cost incurred from Alg. 1. From top to bottom, each row represents target functions $\{g_i\}$ corresponding to those listed in each column of Tab. 1 (from left to right). Shading represents two standard deviations away from the mean Err over 10 independent runs.



Figure 13: Bilevel Alg. 3 applied with K=1 test atoms and 50 gradient descent iterations. Each column visualizes the covariance matrix of the optimized (left) and empirical test distributions (right), respectively. From top to bottom, each row represents target functions $\{g_i\}$ corresponding to those listed in each column of Tab. 1 (from left to right).

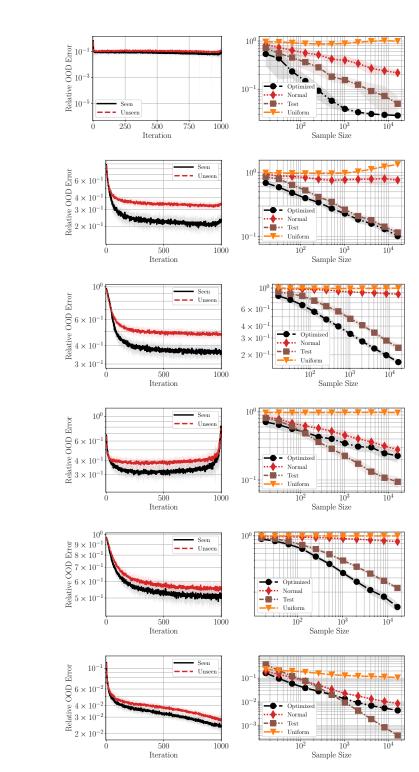


Figure 14: Bilevel Alg. 3 applied to the various ground truths $\{g_i\}$ with K=1 test atoms and 1000 gradient descent iterations. (Left column) Evolution of Err from (F.5) vs. iterations. (Center column) Err of model trained on N samples from optimal ν_{ϑ} vs. nonadaptive distributions. (Right column) Same as center, except incorporating the additional function evaluation cost incurred from Alg. 1. From top to bottom, each row represents target functions $\{g_i\}$ corresponding to those listed in each column of Tab. 1 (from left to right). Shading represents two standard deviations away from the mean Err over 10 independent runs.

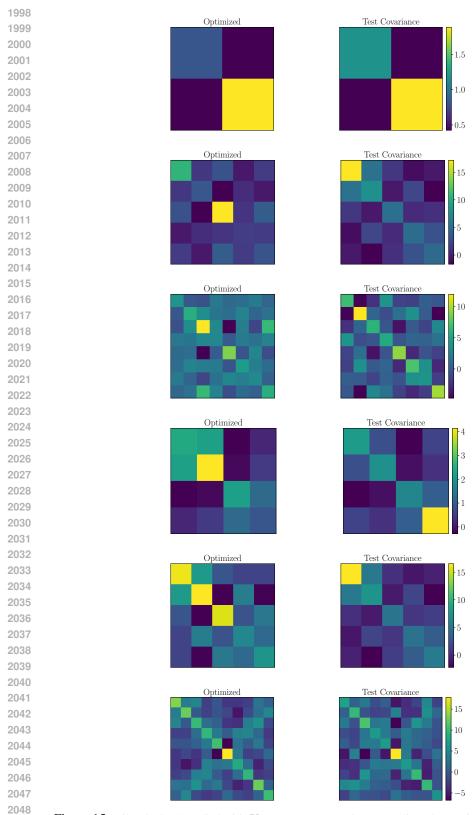


Figure 15: Bilevel Alg. 3 applied with K=1 test atoms and 1000 gradient descent iterations. Each column visualizes the covariance matrix of the optimized (left) and empirical test distributions (right), respectively. From top to bottom, each row represents target functions $\{g_i\}$ corresponding to those listed in each column of Tab. 1 (from left to right).

we must identify an optimal training distribution that enables good generalization across diverse testing scenarios.

To illustrate this, we consider the EIT problem with the AMINO solver; see SM E.3. Suppose the general distribution of interest for the Dirichlet boundary condition is generated according to the law of f_l in (E.4) with $\omega \sim \text{Unif}(\left[\frac{\pi}{2},7\pi\right])$. A reasonable distribution to train upon is this very same distribution. However, suppose we were constrained to focus and train on distributions corresponding to $\text{Law}(\omega)$ supported only on a subset of the domain $\left[\frac{\pi}{2},7\pi\right]$. In that case, the choice of the distribution will greatly affect the OOD error when testing with $\omega \sim \text{Unif}(\left[\frac{\pi}{2},7\pi\right])$. To be more concrete, we will focus on four training distributions for ω , which directly determines the input Dirichlet boundary condition as in SM E.3: $\text{Unif}(\left[\frac{\pi}{2},\pi\right])$, $\text{Unif}(\left[\frac{7\pi}{2},4\pi\right])$, $\text{Unif}(\left[\frac{5\pi}{2},\frac{11\pi}{2}\right])$, and $\text{Unif}(\left[\frac{13\pi}{2},7\pi\right])$. After obtaining the well-trained neural operators (in the sense that the ID generalization error is small enough) corresponding to each training distribution, we study the OOD error of the trained neural operators for Dirichlet data obtained by drawing $\omega \sim \text{Unif}(\left[\frac{\pi}{2},7\pi\right])$ in (E.4).

Model architecture and optimization algorithm. We train four distinct AMINO models; see **SM** E.3 to recall AMINO and **SM** E.1 for its DeepONet component. Each model comprises three jointly trained components:

• Branch encoder: Eight sequential convolutional blocks with LeakyReLU activations:

```
- ConvBlock(1→64, kernel = (1,7), stride = (1,2), pad = (0,3))

- ConvBlock(64→128, kernel = (1,3), stride = (1,2), pad = (0,1))

- ConvBlock(128→128, kernel = (1,3), pad = (0,1))

- ConvBlock(128→256, kernel = (1,3), stride = (1,2), pad = (0,1))

- ConvBlock(256→256, kernel = (1,3), pad = (0,1))

- ConvBlock(256→512, kernel = (1,3), stride = (1,2), pad = (0,1))

- ConvBlock(512→512, kernel = (1,3), pad = (0,1))

- ConvBlock(512→512, kernel = (4,5), pad = 0).
```

The feature map is then flattened and passed through a linear layer to produce the p-dimensional branch output.

- Trunk network: An 8-layer multilayer perceptron (MLP) with 512 neurons per hidden layer, LeakyReLU activations, zero dropout, and a final output dimension of p = 100.
- Fourier Neural Operator (FNO): A single spectral convolution layer with 128 channels and 32 Fourier modes.

All parameters are learned using the AdamW optimizer with a learning rate of 10^{-3} .

Training and optimization procedure. The four neural operator models "Model 1", "Model 2", "Model 3", and "Model 4" are trained using four different input parameter distributions for Law(ω), respectively: Unif($[\frac{\pi}{2},\pi]$), Unif($[\frac{7\pi}{2},4\pi]$), Unif($[5\pi,\frac{11\pi}{2}]$), and Unif($[\frac{13\pi}{2},7\pi]$); see (E.4) to recall how ω is reflected in the actual input distributions on function space. For each distribution, 4096 training samples are generated. Each sample consists of L = 100 Dirichlet–Neumann boundary condition pairs with the underlying conductivity discretized on a 70×70 grid for the domain $\Omega=(0,1)^2$. After 500 training epochs, the trained models are evaluated on 1024 test samples of the distribution for the Dirichlet boundary condition (E.4) determined by $\omega\sim \text{Unif}([\frac{\pi}{2},7\pi])$.

Model evaluation. For both ID and OOD settings, we report the average relative L^1 error given as

Expected Relative
$$L^1$$
 Error $:= \mathbb{E}_{a \sim \text{Law}(a)} \left[\frac{\|a - \mathcal{G}_{\theta} \left((\text{Id}, \Lambda_a)_{\#} \mu \right) \|_{L^1(\Omega)}}{\|a\|_{L^1(\Omega)}} \right],$ (F.6)

where Law(a) is the distribution of a from (E.3) and \mathcal{G}_{θ} is the AMINO model. Using the same setup as in E.3, we approximate (F.6) with a Monte Carlo average over N' = 1024 test conductivity samples $a_n \stackrel{\text{i.i.d.}}{\sim} \text{Law}(a)$ and each with L = 100 pairs of corresponding Dirichlet and Neumann

Table 2: Average relative L^1 error of AMINO models under ID and OOD settings for the EIT inverse problem solver. Model 1 is trained on $\omega_1 \sim \mathsf{Unif}([\frac{\pi}{2},\pi])$, Model 2 on $\omega_2 \sim \mathsf{Unif}([\frac{7\pi}{2},4\pi])$, Model 3 on $\omega_3 \sim \mathsf{Unif}([5\pi,\frac{11\pi}{2}])$, and Model 4 on $\omega_4 \sim \mathsf{Unif}([\frac{13\pi}{2},7\pi])$. In the OOD setting, all models are evaluated on Dirichlet data from $\omega \sim \mathsf{Unif}([\frac{\pi}{2},7\pi])$. One standard deviation is reported. See Fig. 1 for a visual comparison on two representative samples.

	Model 1	Model 2	Model 3	Model 4
In-Distribution (ID)	0.034 ± 0.019	0.081 ± 0.028	0.021 ± 0.010	0.017 ± 0.009
Out-of-Distribution (OOD)	0.796 ± 0.251	0.078 ± 0.047	0.066 ± 0.038	0.132 ± 0.047

boundary data. That is, the expected relative L^1 error is approximately given by the average

$$\frac{1}{N'} \sum_{n=1}^{N'} \frac{\left\| a_n - \mathcal{G}_{\theta} \left(\frac{1}{\mathsf{L}} \sum_{\ell=1}^{\mathsf{L}} \delta_{(f_{\ell}^{(n)}, g_{\ell}^{(a_n)})} \right) \right\|_{L^1}}{\left\| a_n \right\|_{L^1}}, \text{ where } \left\{ (f_{\ell}^{(n)}, g_{\ell}^{(a_n)}) \right\}_{\ell=1}^{\mathsf{L}} \sim \left((\mathrm{Id}, \Lambda_{a_n})_{\#} \mu \right)^{\otimes \mathsf{L}}.$$
(F.7)

The $L^1(\Omega)$ norms in the preceding display are approximated with Riemann sum quadrature over the equally-spaced 70×70 grid.

Tab. 2 illustrates how OOD performance varies with the choice of training distribution. Model 1 achieves strong ID performance but performs poorly in the OOD setting. Model 2 exhibits consistent yet mediocre performance across both ID and OOD evaluations. Model 4 performs well in the ID setting but shows a noticeable drop in OOD performance. Model 3 offers the best balance, achieving low error in both ID and OOD scenarios. Notably, despite being trained on a Dirichlet data distribution with a small support regarding the parameter ω , Model 3 generalizes significantly better to the larger test domain $\left[\frac{\pi}{2}, 7\pi\right]$ for the parameter ω .

These results are further illustrated in Fig. 1, which presents two representative conductivity samples. The top row displays the predicted conductivities from each of the four models when evaluated on Dirichlet–Neumann data pairs sampled from the same distribution used during training. In this ID setting, all four models accurately recover the true conductivity. In contrast, the bottom row shows predictions when the Dirichlet data is drawn from the OOD setting, i.e., $\omega \sim \text{Unif}([\frac{\pi}{2}, 7\pi])$, and the joint Dirichlet–Neumann data distribution also changes accordingly. Under this distributional shift, Models 1, 2, and 4 exhibit noticeable degradation in performance, while Model 3 maintains a reasonable level of accuracy. We use this test to demonstrate the importance of carefully selecting the training data distribution to achieve robust OOD performance.

Computational resources. The AMINO experiment was conducted on a machine equipped with an AMD Ryzen 9 5950X CPU (16 cores / 32 threads), 128 GB of RAM, and an NVIDIA RTX 3080Ti GPU (10 GB VRAM) running with CUDA version 12.1. Twelve CPU and GPU hours total were needed for the training and testing of all four models.

F.3 NTD LINEAR OPERATOR LEARNING

In this example from Sec. 5, we begin by defining a conductivity field $a: \Omega \subset \mathbb{R}^2 \to \mathbb{R}$ given by

$$x \mapsto a(x) = \exp\left(\sum_{j=1}^{M+1} \sum_{k=1}^{M+1} \frac{2\sigma}{\left((j\pi)^2 + (k\pi)^2 + \tau^2\right)^{\alpha/2}} z_{jk} \sin(j\pi x_1) \sin(k\pi x_2)\right)$$
(F.8)

with fixed parameters $\alpha=2$, $\tau=3$, $\sigma=3$ and random coefficients $z_{jk} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0,1)$. Once these random coefficients are drawn, they are fixed throughout the example. Thus, in this experiment, our realization of the conductivity a remains fixed; see the left panel of Fig. 3 for a visualization of this a. Our goal is to develop neural operator approximations of the NtD linear operator $\mathcal{G}^{\star} := \Lambda_a^{-1}$ that maps the Neumann boundary condition to the Dirichlet boundary condition according to the PDE (E.2) (with Neumann data g given instead of Dirichlet data f). Using the alternating minimization algorithm (AMA, Alg. 2) from Sec. 4.2, we want to optimize the training distribution over the Neumann boundary conditions to achieve small OOD error (3.3).

To evaluate the OOD error, we need to assign a random measure \mathbb{Q} . Let $\mathbb{Q} = \frac{1}{K} \sum_{k=1}^{K} \delta_{\nu'_k}$, where $\{\nu'_k\}_{k=1}^{K}$ are distributions over functions. For a fixed k, the function $g \sim \nu'_k$ is a Neumann boundary condition. It takes the form $g \colon \partial \Omega \to \mathbb{R}$, where

$$t \mapsto g(r(t)) = 10\sin(\omega_k(t - \frac{1}{2})) + \sum_{j=1}^{N_{\text{basis}}} \frac{\sigma_k}{\left((j\pi)^2 + \tau_k^2\right)^{\alpha_k/2}} z_j \sqrt{2}\sin(j\pi t). \tag{F.9}$$

We view $r\colon [0,1]\to\partial\Omega$ as a 1D parametrization of the boundary $\partial\Omega\subset\mathbb{R}^2$ for the domain $\Omega:=(0,1)^2$. The entire domain is discretized on an $M\times M$ grid, where M is as in (F.8). Here $z_j\overset{\text{i.i.d.}}{\sim}\mathcal{N}(0,1), \alpha_k=1.5$ for all k, and $N_{\text{basis}}=4M-4$. We set the grid parameter to be M=128 and sample K=64 test distributions in this example. The parameters $\omega_k\overset{\text{i.i.d.}}{\sim}\rho([-2\pi,2\pi])$ and $\tau_k\overset{\text{i.i.d.}}{\sim}\rho([0,50])$ are sampled from a custom distribution $\rho=\rho([x_{\min},x_{\max}])$ and held fixed throughout the algorithm. The custom distribution ρ has the probability density function

$$p(x) = C\sin\left(x - \frac{\pi}{2}\right) + 1\tag{F.10}$$

for $x \in [x_{\min}, x_{\max}] \subset \mathbb{R}$, where C is a normalizing constant ensuring that p integrates to 1 over the domain $[x_{\min}, x_{\max}]$. Samples are generated by using inverse transform sampling. Any time a Neumann boundary condition g is sampled according to (F.9) and (F.10), we subtract the spatial mean of g from itself to ensure that the resulting function integrates to zero. This is necessary to make the Neumann PDE problem—as well as the NtD operator itself—well-posed.

The training distribution ν is assumed to generate functions of the form

$$t \mapsto \widetilde{g}(t) = 10\sin\left(\omega\left(t - \frac{1}{2}\right)\right) + \sum_{i=1}^{N_{\text{basis}}} \frac{\sigma}{\left((j\pi)^2 + \tau^2\right)^{\alpha/2}} \widetilde{z}_j \sqrt{2}\sin(j\pi t), \tag{F.11}$$

where $\widetilde{z}_j \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0,1)$, $\alpha=2,\, \tau=3$, and $\sigma=\tau^{3/2}$. That is, $\nu=\operatorname{Law}(\widetilde{g})$ with \widetilde{g} as in (F.11). The frequency parameter ω appearing in the mean term of (F.11) is optimized using Alg. 2 to identify a training distribution that minimizes average OOD error with respect to \mathbb{Q} .

Note that g and \widetilde{g} are finite-term truncations of the KLE of Gaussian measures with Matérn-like covariance operators; recall Lem. B.2. In particular, for \widetilde{g} , and similarly for g, the covariance operator of the Gaussian measure has eigenpairs

$$\lambda_j = \frac{\sigma^2}{\left((j\pi)^2 + \tau^2\right)^{\alpha}} \quad \text{and} \quad t \mapsto \phi_j(t) = \sqrt{2} \sin(j\pi t) \quad \text{for} \quad j = 1, 2, \dots, N_{\text{basis}}. \tag{F.12}$$

Using (F.12) and (B.1), the squared 2-Wasserstein distance in Prop. 3.1 reduces to

$$W_{2}^{2}(\nu,\nu_{k}') = \int_{0}^{1} \left(10\sin\left(\omega\left(t - \frac{1}{2}\right)\right) - 10\sin\left(\omega_{k}\left(t - \frac{1}{2}\right)\right)\right)^{2} dt + \sum_{j=1}^{N_{\text{basis}}} \left(\sqrt{\lambda_{j}} - \sqrt{\lambda_{j,k}}\right)^{2}, \text{ (F.13)}$$

where $\lambda_{j,k} = \sigma_k^2/((j\pi)^2 + \tau_k^2)^{\alpha_k}$. We approximate the integral with standard quadrature rules. Moreover, the second moment of ν , and similarly for ν_k' , is

$$\mathsf{m}_2(\nu) = \int_0^1 \left(10\sin\left(\omega\left(t - \frac{1}{2}\right)\right)\right)^2 dt + \sum_{j=1}^{N_{\mathrm{basis}}} \lambda_j. \tag{F.14}$$

These closed-form reductions follow directly from the finite-dimensional KLE and Ex. B.1.

Furthermore, we estimate the Lipschitz constants of both the true NtD operator \mathcal{G}^{\star} and the learned model \mathcal{G}_{θ} —which appear in the multiplicative factor c from Prop. 3.1—by sampling 250 realizations of Neumann input function pairs (g_1,g_2) from 11 candidate distributions ν . We denote the finite set of these pairs of realizations by S. We then use the quantity

$$\max_{(g_1, g_2) \in S} \frac{\|\mathcal{G}(g_1) - \mathcal{G}(g_2)\|_{L^2(\partial \Omega)}}{\|g_1 - g_2\|_{L^2(\partial \Omega)}}$$
(F.15)

to estimate the Lipschitz constant of map $\mathcal{G} \in \{\mathcal{G}^*, \mathcal{G}_{\theta}\}$. These Lipschitz constants are then fixed throughout the algorithm.

Model architecture and optimization algorithm. The model architecture is a DeepONet comprising a branch network with five fully connected layers of sizes 4M-4, 512, 512, 512, and 200, and a trunk network with five fully connected layers of sizes 2, 128, 128, 128, and 200, respectively; see **SM** E.1 for details. All layers use ReLU activation functions. Weights are initialized using the Glorot normal initializer. The model is trained with the Adam optimizer with a learning rate of 10^{-3} . After the model-training half step in the alternating minimization scheme Alg. 2, the parameter ω in (F.11) is optimized using SciPy's default implementation of differential evolution with a tolerance of 10^{-7} . Due to the presence of many local minima in the optimization landscape, this global optimization algorithm is necessary.

Training and optimization procedure. The initial value for ω is set to 8π . For every model-training step, 500 pairs of Neumann and Dirichlet data are generated using the current value of ω . For practical reasons, the model is initially trained for 2000 iterations of Adam. If the AMA loss has not improved relative to the previous iteration, training continues for an additional 2000 iterations. This process is repeated up to 10 times, after which the algorithm proceeds to the parameter optimization step for the training distribution ν . During distribution parameter optimization, the AMA loss is evaluated using 500 newly generated data pairs, used specifically to compute the first term of the objective (4.7). To ensure deterministic behavior, the 500 sets of random coefficients $\{z_j\}_{j=1}^{N_{\text{basis}}}$ are fixed separately for training and optimization. This results in 1000 total sets of coefficients fixed for the entire algorithm. This process is repeated across 80 independent runs using the same conductivity but different realizations of the random coefficients $\{z_j\}$. The relative AMA loss and relative OOD error is computed for each run; see (F.16).

Model evaluation. The relative AMA loss is defined as the algorithm's objective value (4.7) normalized by its initial value. The middle panel of Fig. 3 illustrates the rapid decay of this normalized loss over subsequent iterations. The relative OOD loss is defined as

Relative OOD Error :=
$$\sqrt{\frac{\mathbb{E}_{\nu' \sim \mathbb{Q}} \mathbb{E}_{g \sim \nu'} \|\mathcal{G}^* g - \mathcal{G}_{\theta}(g)\|_{L^2}^2}{\mathbb{E}_{\nu' \sim \mathbb{Q}} \mathbb{E}_{g \sim \nu'} \|\mathcal{G}^* g\|_{L^2}^2}},$$
 (F.16)

which measures the squared prediction error of the model \mathcal{G}_{θ} relative to the ground truth NtD operator \mathcal{G}^{\star} on average with respect to the given random measure \mathbb{Q} .

The relative OOD error is estimated using Monte Carlo sampling with K=64 and N'=500 as

Relative OOD Error
$$\approx \sqrt{\frac{\frac{1}{K} \sum_{k=1}^{K} \frac{1}{N} \sum_{n=1}^{N'} \|\mathcal{G}^{\star} g_{k,n} - \mathcal{G}_{\theta}(g_{k,n})\|_{L^{2}}^{2}}{\frac{1}{K} \sum_{k=1}^{K} \frac{1}{N'} \sum_{n=1}^{N'} \|\mathcal{G}^{\star} g_{k,n}\|_{L^{2}}^{2}}},$$
 (F.17)

where $g_{k,n}$ is the n^{th} i.i.d. sample from ν'_k . Further numerical quadrature is used to discretize the $L^2(\partial\Omega)$ norms in the preceding display.

After 80 independent runs of Alg. 2, a 95% confidence interval for the true relative OOD error was computed using a Student's t-distribution as seen in the middle panel of Fig. 3.

Computational resources. The computational resources here are the same as those reported in **SM** F.2, except three total CPU and GPU hours were needed for all independent runs.

F.4 DARCY FLOW FORWARD OPERATOR LEARNING

In this example from Sec. 5, we are interested in learning another operator that is relevant and related to the EIT problem. In contrast with the previous example in **SM** F.3 where we fix the conductivity, the Darcy flow forward operator that we consider is nonlinear and maps the log-conductivity function $\log(a)$ to the PDE solution u of

$$\begin{cases} -\nabla \cdot (a\nabla u) = 1 & \text{in } \Omega, \\ u = 0 & \text{on } \partial\Omega. \end{cases}$$
 (F.18)

That is, \mathcal{G}^* : $\log(a) \mapsto u$. Here, $\Omega := (0,1)^2$.

Again, we first assign a random measure $\mathbb Q$ to evaluate the OOD performance. We take $\mathbb Q=\frac{1}{K}\sum_{k=1}^K \delta_{\nu_k'}$, where each test distribution ν_k' is defined by parameters sampled i.i.d. from custom distributions (F.10): $m_k \sim \rho([0,1])$, $\alpha_k \sim \rho([2,3])$, and $\tau_k \sim \rho([1,4])$. The scaling parameter is set as $\sigma_k = \tau_k^{\alpha_k-1}$. We set K=20 and the grid parameter M=64. Given these parameters, each random log-conductivity function $\log(a')$ from test distribution ν_k' has the form

$$x \mapsto \log(a'(x)) = m_k + \sum_{i=1}^{M+1} \sum_{j=1}^{M+1} \frac{2\sigma_k}{\left((i\pi)^2 + (j\pi)^2 + \tau_k^2\right)^{\alpha_k/2}} z'_{ij} \sin(i\pi x_1) \sin(j\pi x_2), \quad (\text{F.19})$$

where $z'_{ij} \sim \mathcal{N}(0,1)$ independently. Thus, a is lognormal. Sample log-conductivities $\log(a)$ from training distribution ν are assumed to follow the same form

$$x \mapsto \log(a(x)) = m + \sum_{i=1}^{M+1} \sum_{j=1}^{M+1} \frac{2\sigma}{\left((i\pi)^2 + (j\pi)^2 + \tau^2\right)^{\alpha/2}} z_{ij} \sin(i\pi x_1) \sin(j\pi x_2), \quad (F.20)$$

with fixed parameters $\alpha=2$, $\tau=3$, and $\sigma=3$ and $z_{ij} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0,1)$. Alg. 2 is used to optimize the mean shift parameter $m \in \mathbb{R}$ in order to reduce the OOD error. As in **SM** F.3, $\log(a)$ is modeled via a finite KLE of a Gaussian measure with Matérn-like covariance operator. Thus, by working with $\log(a)$ directly, we can exploit (B.1) to write the squared 2-Wasserstein distance factor in objective (4.7) in closed-form.

Model architecture and optimization algorithm. The same DeepONet architecture and training setup from **SM** F.3 are used here, except the branch network has layer sizes M^2 , 256, 256, 128, and 200. After the model-training half step in Alg. 2, the scalar parameter m is optimized using PyTorch's Adam optimizer with a learning rate of 10^{-3} .

Training and optimization procedure. The initial value of m is set to zero. The training and optimization procedure follows the same structure as in SM F.3, but with the following differences: 200 conductivity-solution pairs are generated for training using the current value of m. The model is initially trained for 5000 iterations. If the AMA loss has not improved relative to the previous iteration, training continues for an additional 5000 iterations. This process is repeated up to three times, after which the algorithm proceeds to the distribution parameter optimization step. During distribution parameter optimization, 500 steps are taken. Then the AMA loss is evaluated using three newly generated data pairs. If the AMA loss has not improved relative to the previous iteration, another 500 steps are taken, up to 1500 total steps. After this, the procedure returns to model training again. To ensure deterministic behavior, 203 sets of random coefficients $\{z_{ij}\}$ are fixed across training and optimization. The preceding process is repeated across 13 independent runs, each using different realizations of the coefficients. The average relative OOD error and average relative AMA loss are computed for each run.

Model evaluation. After 13 independent runs, the average relative OOD error and relative AMA loss are computed and are shown in Fig. 16. In Fig. 4, we visualize the pointwise absolute error for iterations 1, 2, 3, and 8. At these iterations, the DeepONet model is trained using the optimal training distribution produced by AMA. The model is then evaluated on the same fixed test conductivity. The progressive reduction in the spatial error across these panels of Fig. 4 corroborates the quantitative decrease in the relative OOD error, which converges to approximately 4% in our final model.

Computational resources. The computational resources here are the same as those reported in **SM** F.3, except three total CPU and GPU hours were needed for all independent runs.

F.5 RADIATIVE TRANSPORT OPERATOR LEARNING

In this final example from Sec. 5, we consider the hyperbolic PDE for radiative transport governing the particle density u(z,v) at position $z \in \Omega \subset \mathbb{R}^d$ and velocity $v \in V \subset \mathbb{R}^d$:

$$\begin{cases} v \cdot \nabla_z u(z,v) = \frac{1}{\varepsilon} a(z) Q[u](z,v), & (z,v) \in \Omega \times V, \\ u(z,v) = \phi(z,v), & (z,v) \in \Gamma_-, \end{cases}$$
 (F.21)

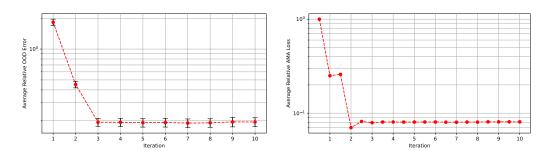


Figure 16: Average relative OOD error and average relative loss of Alg. 2 on the Darcy flow forward problem. (Left) After 13 independent runs, decay of the average relative OOD error of the model when trained on the optimal distribution identified at each iteration; a 95% confidence interval of the true relative OOD error is provided at each iteration. (Right) Decay of average AMA loss defined in (4.7) vs. iteration relative to the same loss at initialization.

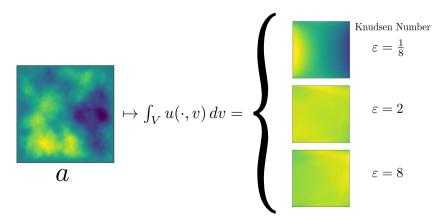


Figure 17: Mapping from scattering coefficient to spatial density in the radiative transport equation for three Knudsen numbers ε . Larger ε values yield a more hyperbolic regime.

where a is the scattering coefficient, ε is the Knudsen number, and the collision operator $u\mapsto Q[u]$ is given by

$$Q[u](z,v) := \int_{V} k(v,v')u(z,v') dv' - u(z,v).$$
 (F.22)

The inflow boundary is defined as

$$\Gamma_{-} := \{ (z, v) \in \partial \Omega \times V \mid -\boldsymbol{n}_z \cdot v \ge 0 \}, \tag{F.23}$$

where n_z denotes the outward normal at $z \in \partial \Omega$.

On three separate experiments with error results given in Fig. 5, we set the Knudsen number ε equal to 1/8, 2, and 8. Changing ε puts us in different physical regimes as seen in Fig. 17. The inflow boundary condition is given by

$$\phi(z,v) = \begin{cases} \left(1 + 0.5\sin(\pi y/L)\right) \max(v \cdot \boldsymbol{n}_z,0), & \text{if } z = 0, \\ 0, & \text{elsewhere on } \Gamma_-. \end{cases} \tag{F.24}$$

We seek the optimal input training distribution ν for learning the nonlinear operator that maps the scattering coefficient function $a(\cdot)$ to the spatial domain density $\int_V u(\cdot,v) \, dv$.

In this example, $\Omega \coloneqq (0,1)^2$ is the unit square and $V \coloneqq \mathbb{S}^1 = \{(\cos\theta,\sin\theta) \mid \theta \in [-\pi,\pi)\}$ is the unit circle. Thus, we parametrize v by θ , i.e., $v(\theta) = (\cos\theta,\sin\theta)$ for $\theta \in [-\pi,\pi)$, $\int_V u(\cdot,v) \, dv = \int_{-\pi}^{\pi} u(\cdot,v(\theta)) \, d\theta$, and $v \cdot \nabla_z = (\cos\theta) \partial_{x_1} + (\sin\theta) \partial_{x_2}$. We first generate samples from \mathbb{Q} to then evaluate the OOD performance of the model trained on ν . We set K=3. Each distribution $\nu_k' \sim \mathbb{Q}$

is set to be $(\exp)_{\#}\gamma_k$, where γ_k is a Gaussian process parametrized by σ_k , α , and τ . Thus, samples from ν'_k have the form

$$a'(x) = \exp\left(f_k(x) + \sum_{i=1}^{\infty} \sum_{j=1}^{\infty} \frac{2\sigma_k}{\left((i\pi)^2 + (j\pi)^2 + \tau^2\right)^{\alpha/2}} z_{ij} \sin(i\pi x_1) \sin(j\pi x_2)\right), \quad (F.25)$$

where $z_{ij} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0,1)$. The mean f_k of each distribution γ_k is selected as follows. For a fixed centered Gaussian process γ , draw $f_k \stackrel{\text{i.i.d.}}{\sim} \gamma$ for each k. Then these f_k functions are then fixed once and for all at the beginning of the experiment. By supplying distinct nonzero means to each of the K test distributions, the problem of finding the optimal training distribution becomes harder. Indeed, the test distributions now exhibit not only covariance shifts, but also mean shifts. We set $\tau=1$ and $\alpha=2$ for all K=3 test distributions, and set $\sigma_1=1$, $\sigma_2=2$, and $\sigma_3=3$. Note that the parameters of γ are set to $\tau=1$, $\alpha=2$, and $\sigma=1$.

Model architecture and optimization algorithm. Following Alg. 2, we represent ν as an empirical measure of N equally weighted particles: $\nu = \frac{1}{N} \sum_{q=1}^{N} \delta_{a_q}$. Each particle a_q is a function of the form

$$a_q(x) = \exp\left(\sum_{i=1}^{\infty} \sum_{j=1}^{\infty} c_{qij} \cdot 2\sin(i\pi x_1)\sin(j\pi x_2)\right),$$
 (F.26)

where the coefficients $c_q := \{c_{qij}\}_{i,j \ge 1}$ belong to ℓ^2 . Rather than optimizing the particles a_q directly, we optimize their representations in the coefficient space through the transformation

$$a_q = \exp(Lc_q), (F.27)$$

where $L\colon\ell^2\to L^2(\Omega)$ is the linear operator mapping coefficients to the series in (F.26) and \exp is applied pointwise. The distribution ν is thus obtained as the pushforward $\nu=(\exp\circ L)_\#\mu$, where μ is the distribution over coefficient sequences. We can then write the loss as $\widetilde{F}(\mu)\coloneqq F((\exp\circ L)_\#\mu)=F(\nu)$, where F is defined as the upper bound in Prop. 3.1, i.e., $\widetilde{F}=F\circ(\exp\circ L)_\#$. For each step k, the discretized Wasserstein gradient flow update rule with respect to $\{c_q\}_{q=1}^N$ is then

$$c_{q}^{(k+1)} = c_{q}^{(k)} - \eta^{(k)} \nabla_{\ell^{2}} \left[D\widetilde{F} \left(\mu_{N}^{(k)} \right) \right] \left(c_{q}^{(k)} \right), \quad \text{where} \quad \mu_{N}^{(k)} \coloneqq \frac{1}{N} \sum_{q'=1}^{N} \delta_{c_{q'}^{(k)}}$$
 (F.28)

and $\eta^{(k)}$ is the step size. By the chain rule and linearity of pushforward,

$$D\widetilde{F}(\mu) = DF((\exp \circ L)_{\#}\mu) \circ (\exp \circ L)$$
 (F.29)

for each μ , where DF denotes the derivative from Prop. 4.9. Thus, we can write the update rule (F.28) in terms of the known derivative DF.

The ℓ^2 gradient with respect to c_q is approximated with PyTorch's autograd. After updating $\{c_q\}_{q=1}^N$, we train a model with discretized versions of particles $\{a_q\}_{q=1}^N$ denoted as $\{\hat{a}_q\}_{q=1}^N$. Each particle a_q is discretized as an $m\times m$ matrix. Similar to the setups from SM F.3 and SM F.4, the model architecture is a DeepONet. However, the branch network now has 6 fully connected layers of sizes m^2 , 100,

Training and optimization procedure. During optimization, we work with N samples from each distribution $\{\nu_k'\}_{k=1}^K$ rather than the distributions themselves. For computational simplicity, we set the number of samples from each ν_k' equal to the number of particles in the training distribution.

The particles $\{a_q\}_{q=1}^N$ are discretized by truncating the number of terms in its expansion, i.e.,

$$a_q(x) = \exp\left(\sum_{i=1}^{N_{\text{truncate}}} \sum_{j=1}^{N_{\text{truncate}}} c_{qij} \cdot 2\sin(i\pi x_1)\sin(j\pi x_2)\right). \tag{F.30}$$

Table 3: Radiative transport operator learning with particle-based Alg. 2. Reported are mean relative OOD errors with 95% confidence intervals computed over 10 independent runs for three separate models with varying Knudsen number ε . All models have the same DeepONet architecture and are trained with N=336 samples for 5,000 epochs. What makes each model different is their training distribution. The initial model was trained with random particles $c_{qij} \sim \mathcal{N}(0, 10^{-6})$, the optimized model was trained with particles after one iteration of particle-based AMA, and the benchmark model is a model trained with samples from the mixture of test distributions, i.e., $\nu_{\mathbb{Q}} = \frac{1}{K} \sum_{k=1}^{K} \nu_k'$.

ε	Initial	Optimized	Benchmark
1/8	$\begin{array}{ c c c c c }\hline 8.93 \times 10^{-2} \pm 5.87 \times 10^{-3} \\ 9.12 \times 10^{-2} \pm 6.28 \times 10^{-3}\end{array}$	$1.36 \times 10^{-2} \pm 1.54 \times 10^{-3}$	$1.09 \times 10^{-2} \pm 4.22 \times 10^{-4}$
$2^{'}$	$9.12 \times 10^{-2} \pm 6.28 \times 10^{-3}$	$1.16 \times 10^{-2} \pm 2.00 \times 10^{-3}$	$9.26 \times 10^{-3} \pm 2.01 \times 10^{-3}$
8	$9.04 \times 10^{-2} \pm 6.75 \times 10^{-3}$	$1.16 \times 10^{-2} \pm 3.84 \times 10^{-3}$	$8.71 \times 10^{-3} \pm 1.80 \times 10^{-3}$

The coefficient sequences $\{c_q\}_{q=1}^N$ are truncated to $N_{\text{truncate}}=10$ for both indices i and j, yielding sequences of 100 coefficients each. These are initialized by sampling from $\mathcal{N}(0, 10^{-6})$.

We discretize Ω on a 51×51 grid (m=51) and V with 12 different directions, i.e. $\Delta\theta=\pi/6$, so we approximate the spatial-domain density by

$$\tilde{u}_q \coloneqq \int_V u_q(\cdot, v) \, dv \approx \sum_{i=1}^{12} u_q(\cdot, \theta_i) \, \Delta\theta \,. \tag{F.31}$$

At each iteration of Alg. 2:

- 1. Compute the true outputs $\{\tilde{u}_q\}_{q=1}^N$ as $m \times m$ matrices using the current $\{c_q\}_{q=1}^N$.
- 2. Train the DeepONet for 5,000 epochs with an 80/20 train-test split.
- 3. Update $\{c_q\}_{q=1}^N$ according to Eqn. (F.28) with initial step size $\eta=0.0125$.
- 4. If the loss $F(\nu)$ increases, halve the step size and recompute the update.

In this example, only one iteration is performed because convergence occurs quickly and no more significant improvement is made in later iterations.

For the loss computation in Step 4 above, the second moment of a distribution ν is estimated with N samples by

$$\mathsf{m}_{2}(\nu) = \mathbb{E}_{a \sim \nu} \|a\|_{L^{2}(\Omega)}^{2} \approx \frac{1}{N} \sum_{i=1}^{N} \|a_{i}\|_{L^{2}(\Omega)}^{2},$$
 (F.32)

where $||a_i||_{L^2(\Omega)} = \int_{\Omega} a_i(z)^2 dz$ is estimated with the trapezoidal rule. Note that for our empirical training distribution, F.32 is not an approximation, but instead an equality. We estimate the Wasserstein distance $W_2(\nu,\nu')$ between distributions ν and ν' using N samples from each, with the ground cost function taken to be $c(a,a') = ||a-a'||_{L^2(\Omega)}^2$.

Model evaluation. We compare models trained on the optimized distribution at each iteration of particle-based AMA with a benchmark model trained on N samples from the mixture of test distributions $\nu_{\mathbb{Q}} = \frac{1}{K} \sum_{k=1}^{K} \nu_k'$ (Eqn. D.1), using the same DeepONet architecture and 5,000 training epochs for all models. Performance is measured by the relative OOD error (Eqn. F.16), approximated with the original N samples from each ν_k' . As shown in Fig. 5, the optimized model improves substantially after just one iteration, achieving performance comparable to the benchmark. Given that AMA minimizes an upper bound, we are not directly minimizing the OOD error which can account for why the optimized distribution does not outperform the benchmark distribution for most sample sizes. Tab. 3 reports 95% confidence intervals for relative OOD errors over 10 independent trials at N=336, demonstrating that the optimized model reduces its relative OOD error by 85% from its original value on average after one iteration. For this particular sample size, the runtime per iteration, shown in Fig. 18, averages 95 seconds for $\varepsilon=\frac{1}{8}$, 165 seconds for $\varepsilon=2$, and 224 seconds for $\varepsilon=8$.

Computational resources. The computational resources here are the same as those reported in **SM** F.3, except 18 total CPU and GPU hours were needed for all independent runs.

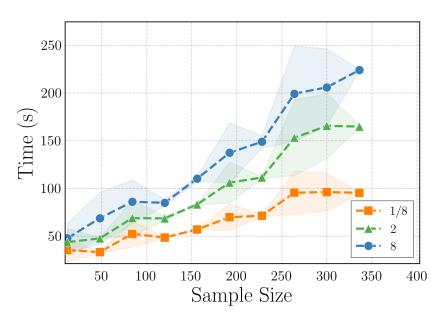


Figure 18: Runtime of one particle-based AMA iteration for the radiative transport equation with sample size N at three Knudsen numbers $\varepsilon \in \{1/8,2,8\}$. Error bars indicate 95% confidence intervals over 10 trials.