

LEVERAGING MULTITIME HAMILTON–JACOBI PDEs FOR CERTAIN SCIENTIFIC MACHINE LEARNING PROBLEMS*

PAULA CHEN[†], TINGWEI MENG[‡], ZONGREN ZOU[†], JÉRÔME DARBON[§],
AND GEORGE EM KARNIADAKIS[¶]

Abstract. Hamilton–Jacobi partial differential equations (HJ PDEs) have deep connections with a wide range of fields, including optimal control, differential games, and imaging sciences. By considering the time variable to be a higher dimensional quantity, HJ PDEs can be extended to the multitime case. In this paper, we establish a novel theoretical connection between specific optimization problems arising in machine learning and the multitime Hopf formula, which corresponds to a representation of the solution to certain multitime HJ PDEs. Through this connection, we increase the interpretability of the training process of certain machine learning applications by showing that when we solve these learning problems, we also solve a multitime HJ PDE and, by extension, its corresponding optimal control problem. As a first exploration of this connection, we develop the relation between the regularized linear regression problem and the linear quadratic regulator (LQR). We then leverage our theoretical connection to adapt standard LQR solvers (namely, those based on the Riccati ordinary differential equations) to design new training approaches for machine learning. Finally, we provide some numerical examples that demonstrate the versatility and possible computational advantages of our Riccati-based approach in the context of continual learning, posttraining calibration, transfer learning, and sparse dynamics identification.

Key words. multitime Hamilton–Jacobi PDEs, Hopf formula, machine learning, linear quadratic regulator, linear regression, Riccati equation

MSC codes. 35F21, 49N05, 49N10, 68T05, 35B37

DOI. 10.1137/23M1561397

1. Introduction. It is well known that Hamilton–Jacobi partial differential equations (HJ PDEs) have deep connections to optimal control [3], differential games [19], and imaging sciences [12, 16], among many other fields. When the Hamiltonians are convex and only depend on the momentum, the solution to the HJ PDEs can be represented by a Hopf formula, which converts the solution of the PDE to the solution of an optimization problem. Multitime HJ PDEs were originally introduced in economics [42]. The solution to certain multitime HJ PDEs was then shown to be able to be represented by a multitime Hopf formula [33], which is a generalization of the single-time case and has been shown to have connections with imaging sciences [15].

In this paper, we establish a novel theoretical connection between certain optimization problems arising in machine learning and the multitime Hopf formula

*Submitted to the journal’s Machine Learning Methods for Scientific Computing section March 24, 2023; accepted for publication (in revised form) November 6, 2023; published electronically March 15, 2024. Paula Chen, Tingwei Meng, and Zongren Zou contributed equally to this work.

<https://doi.org/10.1137/23M1561397>

Funding: The first author is supported by the SMART Scholarship, which is funded by the Under Secretary of Defense/Research and Engineering (USD/R&E), National Defense Education Program (NDEP)/ BA-1, Basic Research. The third, fourth and fifth authors are supported by the MURI/AFOSR FA9550-20-1-0358 project. We also acknowledge the support by award DOE-MMICS SEA-CROGS DE-SC0023191.

[†]Division of Applied Mathematics, Brown University, Providence, RI 02912 USA (paula.chen@alumni.brown.edu, zongren.zou@brown.edu).

[‡]Department of Mathematics, UCLA, Los Angeles, CA 90025 USA (tingwei@math.ucla.edu).

[§]Corresponding author. Division of Applied Mathematics, Brown University, Providence, RI 02912 USA (jerome.darbon@brown.edu).

[¶]Division of Applied Mathematics, Brown University, Providence, RI 02912 USA, and Pacific Northwest National Laboratory, Richland, WA 99354 USA (george.karniadakis@brown.edu).

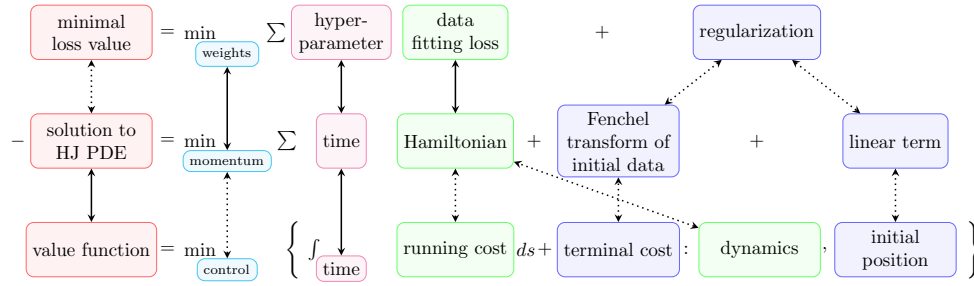


FIG. 1. (See section 2). Illustration of a connection between a regularized learning problem (top), the multitime Hopf formula (middle), and an optimal control problem (bottom). The colors indicate the associated quantities between each problem. For example, the optimal weights in the learning problem are equivalent to the momentum in the HJ PDE, which is related to the control in the optimal control problem (cyan), and the hyperparameters for the data fitting terms correspond to the time variables in the HJ PDE and the time horizon in the optimal control problem (magenta). This color scheme is reused in the subsequent illustrations of our connection. The solid-line arrows denote direct equivalences. The dotted arrows represent additional mathematical relations. Note: color appears only in the online article.

(section 2). Specifically, we show that there are one-to-one correspondences between the loss functions in regularized learning problems and the objective function of the multitime Hopf formula, which in turn yields connections to optimal control. See Figure 1 for an illustration of these correspondences. As such, our connection increases the interpretability of the training process of certain machine learning applications by showing that when we solve these learning problems, we actually solve a multitime HJ PDE and, by extension, its corresponding optimal control problem. In this paper, we show that our connection allows us to leverage HJ PDEs and optimal control theory and algorithms to solve optimization problems arising from machine learning applications, and we reserve the reverse direction for future work.

As a first exploration of this connection, we focus on the connection between regularized linear regression problems and the Linear Quadratic Regulator (LQR) (section 3). We use this case to gain a deeper understanding of our connection and illustrate its potential benefits. Linear regression consists of learning a linear prediction model and is one of the fundamental learning problems in supervised machine learning [43, 37, 49]. LQR [46, 2] is a well-studied optimal control problem with quadratic running and terminal costs and linear dynamics and is typically solved using the Riccati ordinary differential equations (ODEs) [36, 13, 38]. Through our connection, we establish that solving regularized linear regression problems is equivalent to solving particular LQR problems. As a result, we can leverage standard LQR solvers to design new training approaches for machine learning. In particular, we develop new methodology for solving regularized linear regression problems by adapting solvers for the Riccati ODEs to these new settings (section 4).

To highlight the versatility and possible computational advantages of our new Riccati-based approach, we apply our methodology to several test problems in machine learning (section 5). In the first example, we consider a function approximation problem to demonstrate the computational and memory advantages of our Riccati-based approach in the context of continual learning [39, 28, 47]. Specifically, we show that our Riccati-based approach naturally enables us to continually adapt the learned model to new data without having to store or retrain on the previous data (which is especially significant given the rise of big data), while also avoiding catastrophic

forgetting [28, 39]. In the second example, we demonstrate how our Riccati-based approach can be used to perform posttraining calibrations. In particular, our approach gives us the flexibility to add or remove data points and tune hyperparameters to increase the accuracy of the learned model without having to retrain it entirely, which again provides computational and memory advantages over conventional learning methods. In the third example, we use our Riccati-based methodology to fit the last layer of a physics-informed neural network (PINN) [41] using transfer learning [52, 18]. In this application, we demonstrate that as we change the value of the hyperparameters, solving the associated Riccati ODEs not only provides the solution to the updated problem, but also a *continuum* of solutions along a one-dimensional (1D) curve on the Pareto front of the data fitting losses and regularization. Finally, in the fourth example, we highlight the versatility of our Riccati-based approach by showing how it can be combined with existing optimization methods (e.g., the primal-dual hybrid gradient (PDHG) algorithm [7]) to perform sparse dynamics identification [5].

The main contributions of this work are the development of a new theoretical connection between learning problems and the Hopf formula and a new Riccati-based approach for solving regularized linear regression problems. While we demonstrate promising results, the work presented here has some limitations. For example, while we establish our theoretical connection between more general learning problems, multitime HJ PDEs, and optimal control problems, we have yet to fully explore nonlinear learning models, general convex Hamiltonians, or nonlinear control dynamics. Additionally, as discussed previously, we have also not yet investigated what possible advantages our connection provides for solving HJ PDEs and optimal control problems. In particular, many efficient solvers for high-dimensional problems in machine learning exist [29, 20]; it would be desirable to be able to leverage our connection to reuse this machinery for HJ PDEs and optimal control. Thus, our novel connection presents many exciting opportunities. We discuss some other possible future directions in section 6. The overall structure of this paper is illustrated in Figure 2.

2. Generalized Hopf formula. In this section, we provide some mathematical background on the single-time and multitime Hopf formulas. Specifically, we review the well-known connections between the Hopf formula, the solution to HJ PDEs with time- and space-independent Hamiltonians, and the solution to the corresponding optimal control problems. We then present a novel theoretical connection between the Hopf formula and regularized learning problems. Through this connection, we establish that when we solve these learning problems, we actually evaluate the solution to certain HJ PDEs and their corresponding optimal control problems, and vice versa.

2.1. Introduction to the Hopf formula. The single-time HJ PDE is

$$(2.1) \quad \begin{cases} \frac{\partial S(\mathbf{x}, t)}{\partial t} + H(\nabla_{\mathbf{x}} S(\mathbf{x}, t)) = 0, & \mathbf{x} \in \mathbb{R}^n, t > 0, \\ S(\mathbf{x}, 0) = J(\mathbf{x}), & \mathbf{x} \in \mathbb{R}^n, \end{cases}$$

where $H: \mathbb{R}^n \rightarrow \mathbb{R}$ is the Hamiltonian and $J: \mathbb{R}^n \rightarrow \mathbb{R}$ is the initial condition. Assume that H and J are convex (although we note that these assumptions could be relaxed; e.g., see [24]). Then, the viscosity solution to the single-time HJ PDE (2.1) is given by the Hopf formula [24]

$$(2.2) \quad S(\mathbf{x}, t) = \sup_{\mathbf{p} \in \mathbb{R}^n} \{ \langle \mathbf{x}, \mathbf{p} \rangle - tH(\mathbf{p}) - J^*(\mathbf{p}) \} = - \inf_{\mathbf{p} \in \mathbb{R}^n} \{ tH(\mathbf{p}) + J^*(\mathbf{p}) - \langle \mathbf{x}, \mathbf{p} \rangle \},$$

where f^* denotes the Fenchel–Legendre transform of the function f ; i.e., $f^*(p) = \sup_{x \in \mathbb{R}^n} \{ \langle x, p \rangle - f(x) \}$.

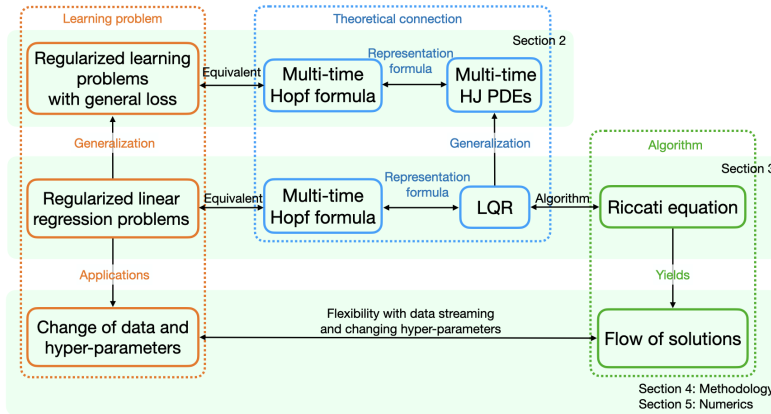


FIG. 2. Overview of the overall structure of this paper, the learning problems considered, our new theoretical connection, and our new Riccati-based algorithm. We show the learning problem in the left (orange), the theoretical connection in the middle (blue), and the algorithm in the right (green). The top row contains the general connection between learning problems and HJ PDEs (see section 2), the middle row contains the connection between the regularized linear regression problem and LQR (see section 3), and the bottom row contains the connection between applications in machine learning and our Riccati-based algorithm (see sections 4 and 5). Note: color appears only in the online article.

The value function of the following optimal control problem also solves (2.1):

$$(2.3) \quad S(\mathbf{x}, t) = \min_{\mathbf{u}(\cdot)} \left\{ \int_0^t L(\mathbf{u}(s)) ds + J(\mathbf{x}(t)) : \dot{\mathbf{x}}(s) = f(\mathbf{u}(s)) \forall s \in (0, t], \mathbf{x}(0) = \mathbf{x} \right\},$$

where the running cost L and the source term f of the dynamics are related to the Hamiltonian H by $H(\mathbf{p}) = \sup_{\mathbf{u} \in \mathbb{R}^m} \{ \langle -f(\mathbf{u}), \mathbf{p} \rangle - L(\mathbf{u}) \}$ and, in this context, we interpret J to be the terminal cost.

A natural generalization of this formulation to the multitime case is as follows. Let H_1, \dots, H_N be convex Hamiltonians, such that $\text{dom } H_i = \mathbb{R}^n$ for all $i = 1, \dots, N$, and let J be a convex initial condition. (Again, we note that these assumptions could be relaxed; e.g., see [33].) Then, the multitime HJ PDE is given by

$$(2.4) \quad \begin{cases} \frac{\partial S(\mathbf{x}, t_1, \dots, t_N)}{\partial t_i} + H_i(\nabla_{\mathbf{x}} S(\mathbf{x}, t_1, \dots, t_N)) = 0 \text{ for } i \in \{1, \dots, N\}, & \mathbf{x} \in \mathbb{R}^n, t_1, \dots, t_N > 0, \\ S(\mathbf{x}, 0, \dots, 0) = J(\mathbf{x}), & \mathbf{x} \in \mathbb{R}^n, \end{cases}$$

and the solution to the multitime HJ PDE (2.4) can be represented by the following generalized (multitime) Hopf formula [33]:

$$(2.5) \quad \begin{aligned} S(\mathbf{x}, t_1, \dots, t_N) &= \sup_{\mathbf{p} \in \mathbb{R}^n} \left\{ \langle \mathbf{x}, \mathbf{p} \rangle - \sum_{i=1}^N t_i H_i(\mathbf{p}) - J^*(\mathbf{p}) \right\} \\ &= - \inf_{\mathbf{p} \in \mathbb{R}^n} \left\{ \sum_{i=1}^N t_i H_i(\mathbf{p}) + J^*(\mathbf{p}) - \langle \mathbf{x}, \mathbf{p} \rangle \right\}. \end{aligned}$$

Moreover, the value function of an optimal control problem in the form of (2.3) with terminal time $t = \sum_{j=1}^N t_j$ and where the running cost L and the source term




FIG. 3. (See section 2). Mathematical formulation describing the connection between a regularized learning problem (top), the multitime Hopf formula (middle), and an optimal control problem (bottom). The content of this illustration matches that of Figure 1 by replacing each term in Figure 1 with its corresponding mathematical expression. The colors indicate the associated quantities between each problem. The solid-line arrows denote direct equivalences. The dotted arrows represent additional mathematical relations. Note: color appears only in the online article.

f of the dynamics are defined piecewise by $L(s, \mathbf{u}) = L_i(\mathbf{u})$ and $f(s, \mathbf{u}) = f_i(\mathbf{u})$, respectively, for $s \in \left(\sum_{j=1}^{i-1} t_j, \sum_{j=1}^i t_j\right]$ and $i = 1, \dots, N$, also solves the HJ PDE (2.4). The piecewise running costs L_i and piecewise source terms f_i of the dynamics are related to the Hamiltonians H_i by $H_i(\mathbf{p}) = \sup_{\mathbf{u} \in \mathbb{R}^m} \{-f_i(\mathbf{u}, \mathbf{p}) - L_i(\mathbf{u})\}$.

2.2. Connection between the Hopf formula and learning problems.

In this section, we connect the Hopf formulas (2.2) and (2.5) with regularized learning problems. Consider a learning problem with data points $\{(\mathbf{z}_i, \mathbf{y}_i)\}_{i=1}^N \subset \mathbb{R}^M \times \mathbb{R}^m$. The goal of the learning problem is to find a function $F(\mathbf{z}; \boldsymbol{\theta})$ with input $\mathbf{z} \in \mathbb{R}^M$ and unknown parameter $\boldsymbol{\theta} \in \mathbb{R}^n$, such that $\mathcal{A}F(\mathbf{z}_i; \boldsymbol{\theta})$ approximately equals \mathbf{y}_i at every \mathbf{z}_i . Here \mathcal{A} is an operator acting on the function F . For instance, \mathcal{A} could be the identity operator (as in regression problems [49]) or a differential operator (as in PINNs [41]). Then, the learning problem is given by the following optimization problem:

$$(2.6) \quad \min_{\boldsymbol{\theta} \in \mathbb{R}^n} \sum_{i=1}^N \lambda_i \mathcal{L}_i(\mathcal{A}F(\mathbf{z}_i; \boldsymbol{\theta}), \mathbf{y}_i) + R(\boldsymbol{\theta}).$$

The above loss function consists of two parts: each $\mathcal{L}_i(\mathcal{A}F(\mathbf{z}_i; \boldsymbol{\theta}), \mathbf{y}_i)$ is a data fitting term at $(\mathbf{z}_i, \mathbf{y}_i)$ (where $\mathcal{L}_i(\mathbf{a}, \mathbf{b})$ is a function measuring the discrepancy between \mathbf{a} and \mathbf{b}) and R is a regularization term. In this paper, we assume the functions $\boldsymbol{\theta} \mapsto \mathcal{L}_i(\mathcal{A}F(\mathbf{z}_i; \boldsymbol{\theta}), \mathbf{y}_i)$ are convex for all $i = 1, \dots, N$.

Then, the connection between the learning problem (2.6), the Hopf formulas (2.2) and (2.5), and the optimal control problem (2.3) is illustrated in Figure 3. Specifically, if there is only one data point ($N = 1$), the learning problem (2.6) is related to the single-time Hopf formula (2.2) by setting $\boldsymbol{\theta} = \mathbf{p}$, $H(\mathbf{p}) = \mathcal{L}_1(\mathcal{A}F(\mathbf{z}_1; \mathbf{p}), \mathbf{y}_1)$, $t = \lambda_1$, and $R(\mathbf{p}) = J^*(\mathbf{p}) - \langle \mathbf{x}, \mathbf{p} \rangle + c(\mathbf{x})$, where $c(\mathbf{x})$ is a constant (possibly 0) that is independent of \mathbf{p} but may depend on \mathbf{x} . In other words, the variables \mathbf{x}, t in the HJ PDE become hyperparameters in the learning problem, and we can treat them as constants when optimizing the learning problem (2.6) with respect to $\boldsymbol{\theta} = \mathbf{p}$. Hence, when we solve these single-point learning problems, we simultaneously evaluate the solution to the HJ PDE (2.1) at the point (\mathbf{x}, t) or, equivalently, we solve the corresponding optimal control problem (2.3). Conversely, when we solve the HJ PDE (2.1), the spatial gradient $\nabla_{\mathbf{x}} S(\mathbf{x}, t)$ of the solution gives the minimizer $\boldsymbol{\theta}^*$ to the single-point learning problem (2.6).

If there are $N > 1$ data points, then the learning problem (2.6) is related to the multitime Hopf formula (2.5) by setting $\boldsymbol{\theta} = \mathbf{p}$, $H_i(\mathbf{p}) = \mathcal{L}_i(\mathcal{A}F(\mathbf{z}_i; \mathbf{p}), \mathbf{y}_i)$, $t_i = \lambda_i$, and $R(\mathbf{p}) = J^*(\mathbf{p}) - \langle \mathbf{x}, \mathbf{p} \rangle + c(\mathbf{x})$, where $c(\mathbf{x})$ is a constant (possibly 0) that is independent of \mathbf{p} but may depend on \mathbf{x} . Hence, solving these multipoint learning problems is equivalent to evaluating the solution to the HJ PDE (2.5) at the point $(\mathbf{x}, t_1, \dots, t_N)$, which is also equivalent to solving the corresponding optimal control problem (2.3) with terminal time $t = \sum_{j=1}^N t_j$ and piecewise running costs L_i and dynamics f_i related via $H_i(\mathbf{p}) = \sup_{\mathbf{u} \in \mathbb{R}^m} \{ \langle -f_i(\mathbf{u}), \mathbf{p} \rangle - L_i(\mathbf{u}) \}$. Similarly to the single-time case, when we solve the multitime HJ PDE (2.4), the spatial gradient $\nabla_{\mathbf{x}} S(\mathbf{x}, t_1, \dots, t_N)$ of the solution gives the minimizer $\boldsymbol{\theta}^*$ to the multipoint learning problem (2.6). We provide a glossary of the mathematical relationships between these three problem formulations in Appendix A.

3. LQR. In this section, we develop our theoretical connection from section 2.2 in the specific case where the optimal control problem (2.3) corresponds to the LQR problem [46, 2]. We show that solving certain LQR problems is equivalent to solving learning problems with linear models, quadratic data fitting losses, and quadratic regularization (i.e., an ℓ_2 -regularized linear regression problem). Although broader classes of learning problems are of interest, we restrict ourselves to linear regression problems as a starting point for demonstrating the potential of our theoretical connection. Specifically, we leverage our theoretical connection to show how established techniques for solving HJ PDEs (e.g., the Riccati ODEs [36]) can be reused to solve this class of learning problems.

3.1. Introduction to the LQR and Riccati equation. The finite-horizon, continuous-time LQR is given by

$$(3.1) \quad S(\mathbf{x}, t) = \min_{\mathbf{u}(\cdot)} \left\{ \int_0^t \left(\frac{1}{2} \mathbf{x}(s)^T \mathcal{Q} \mathbf{x}(s) + \frac{1}{2} \mathbf{u}(s)^T \mathcal{R} \mathbf{u}(s) + \mathbf{x}(s)^T \mathcal{N} \mathbf{u}(s) \right) ds + \frac{1}{2} \mathbf{x}(t)^T \mathcal{Q}_f \mathbf{x}(t) : \dot{\mathbf{x}}(s) = \mathbf{A} \mathbf{x}(s) + \mathbf{B} \mathbf{u}(s) \forall s \in (0, t], \mathbf{x}(0) = \mathbf{x} \right\},$$

where $\mathcal{Q}, \mathcal{Q}_f \in \mathbb{R}^{n \times n}$ and $\mathcal{R} \in \mathbb{R}^{m \times m}$ are symmetric positive definite, $\mathcal{N} \in \mathbb{R}^{n \times m}$, $\mathbf{A} \in \mathbb{R}^{n \times n}$, and $\mathbf{B} \in \mathbb{R}^{n \times m}$. The corresponding HJ PDE is

$$(3.2) \quad \begin{cases} \frac{\partial S(\mathbf{x}, t)}{\partial t} + H(\mathbf{x}, \nabla_{\mathbf{x}} S(\mathbf{x}, t)) = 0, & \mathbf{x} \in \mathbb{R}^n, t > 0, \\ S(\mathbf{x}, 0) = J(\mathbf{x}), & \mathbf{x} \in \mathbb{R}^n, \end{cases}$$

where the initial data of the HJ PDE are given by the terminal cost $J(\mathbf{x}) := \frac{1}{2} \mathbf{x}^T \mathcal{Q}_f \mathbf{x}$ of the optimal control problem and the Hamiltonian H is defined by

$$(3.3) \quad \begin{aligned} H(\mathbf{x}, \mathbf{p}) &= \sup_{\mathbf{u} \in \mathbb{R}^m} \langle -f(\mathbf{x}, \mathbf{u}), \mathbf{p} \rangle - L(\mathbf{x}, \mathbf{u}) \\ &= -\langle \mathbf{A} \mathbf{x}, \mathbf{p} \rangle - \frac{1}{2} \langle \mathbf{x}, \mathcal{Q} \mathbf{x} \rangle + \frac{1}{2} \langle \mathbf{B}^T \mathbf{p} + \mathcal{N}^T \mathbf{x}, \mathcal{R}^{-1} (\mathbf{B}^T \mathbf{p} + \mathcal{N}^T \mathbf{x}) \rangle, \end{aligned}$$

where $f(\mathbf{x}, \mathbf{u}) = \mathbf{A} \mathbf{x} + \mathbf{B} \mathbf{u}$ is the source term of the dynamics and $L(\mathbf{x}, \mathbf{u}) = \frac{1}{2} \mathbf{x}^T \mathcal{Q} \mathbf{x} + \frac{1}{2} \mathbf{u}^T \mathcal{R} \mathbf{u} + \mathbf{x}^T \mathcal{N} \mathbf{u}$ is the running cost. Note that because of the spatial dependence in the Hamiltonian H , the Hopf formula cannot be applied directly to the above LQR problem without additional assumptions. In sections 3.2 and 3.3, we discuss some assumptions under which H is independent of the spatial variable \mathbf{x} and the




FIG. 4. (See section 3). Mathematical formulation describing the connection between a regularized linear regression problem (top), the multitime Hopf formula (middle), and a piecewise LQR problem with $A = Q = 0$ and $\mathcal{N} = 0$ on each piece (bottom). Note that $T_i = \sum_{j=1}^i t_j$. The content of this illustration is a special case of the connection in Figure 3 using quadratic data fitting losses and quadratic regularization. The colors indicate the associated quantities between each problem. The solid-line arrows denote direct equivalences. The dotted arrows represent additional mathematical relations. Note: color appears only in the online article.

corresponding learning problems that can be solved via our connection through the Hopf formula (see, for example, Figure 4).

It is well-known that this LQR problem can be solved using the Riccati equation as follows [36]. Define $C_{pp} = B\mathcal{R}^{-1}B^T$, $C_{xx} = -\mathcal{N}\mathcal{R}^{-1}\mathcal{N}^T + Q$, and $C_{xp} = A - B\mathcal{R}^{-1}\mathcal{N}^T$. Then, the solution is also given by $S(\mathbf{x}, t) = \frac{1}{2}\mathbf{x}^T P(t)\mathbf{x}$, where the function $P : [0, \infty) \rightarrow \mathbb{R}^{n \times n}$ takes values in the space of symmetric positive definite matrices and solves the following Riccati equation:

$$(3.4) \quad \begin{cases} \dot{P}(t) = -P(t)^T C_{pp} P(t) + P(t)^T C_{xp} + C_{xp}^T P(t) + C_{xx} & t \in (0, +\infty), \\ P(0) = Q_f. \end{cases}$$

When H does not depend on \mathbf{x} , we can use our connection to modify the corresponding LQR problem to consider different HJ PDEs and, hence, to accommodate different learning problems. For example, adding lower order terms in the running cost L and/or the source term f of the dynamics corresponds to adding lower order terms in the Hamiltonian of the HJ PDE or, equivalently, in the data fitting term of the learning problem. Similarly, adding lower order terms in the terminal cost J corresponds to adding lower order terms in the initial condition of the HJ PDE or, equivalently, in the regularization term of the learning problem. In Appendix B, we present a more general LQR problem with lower order terms that is more amenable to forming our connection to linear regression problems.

3.2. Connection to single-point regularized linear regression problems.

In this section, we establish a relation between the linear regression problem with only one data point and the LQR problem (3.1) with $A = Q = 0$ and $\mathcal{N} = 0$. Note that, to do this, we also add in some lower order terms to the original LQR problem (3.1) (e.g., see Appendix B). In this case, the LQR problem becomes

$$(3.5) \quad S(\mathbf{x}, t) = \min_{\mathbf{u}(\cdot)} \left\{ \int_0^t \left(\frac{1}{2} \mathbf{u}(s)^T \mathcal{R} \mathbf{u}(s) - \mathbf{a}^T \mathbf{u}(s) \right) ds + \frac{1}{2} \mathbf{x}(t)^T Q_f \mathbf{x}(t) + \mathbf{b}^T \mathbf{x}(t) : \dot{\mathbf{x}}(s) = B \mathbf{u}(s) \forall s \in (0, t], \mathbf{x}(0) = \mathbf{x} \right\},$$

and the corresponding HJ PDE is given by (3.2), where $J(\mathbf{x}) = \frac{1}{2}\mathbf{x}^T \mathcal{Q}_f \mathbf{x} + \mathbf{b}^T \mathbf{x}$ is the initial data/terminal cost whose Fenchel transform is given by

$$J^*(\mathbf{p}) = \sup_{\mathbf{x} \in \mathbb{R}^n} \langle \mathbf{x}, \mathbf{p} \rangle - J(\mathbf{x}) = \frac{1}{2} \left\| \mathcal{Q}_f^{-1/2}(\mathbf{p} - \mathbf{b}) \right\|_2^2$$

and the Hamiltonian H is given by

$$H(\mathbf{p}) = \sup_{\mathbf{u} \in \mathbb{R}^m} \langle -B\mathbf{u}, \mathbf{p} \rangle - \frac{1}{2} \mathbf{u}^T \mathcal{R} \mathbf{u} + \mathbf{a}^T \mathbf{u} = \frac{1}{2} \left\| \mathcal{R}^{-1/2}(B^T \mathbf{p} - \mathbf{a}) \right\|_2^2,$$

where $f(\mathbf{u}) = B\mathbf{u}$ is the source term of the dynamics and $L(\mathbf{u}) = \frac{1}{2} \mathbf{u}^T \mathcal{R} \mathbf{u}$ is the running cost. Then, using the single-time Hopf formula (2.2), we have that the solution to the HJ PDE is given by

$$(3.6) \quad S(\mathbf{x}, t) = \sup_{\mathbf{p} \in \mathbb{R}^n} \left\{ \langle \mathbf{x}, \mathbf{p} \rangle - \frac{t}{2} \left\| \mathcal{R}^{-1/2}(B^T \mathbf{p} - \mathbf{a}) \right\|_2^2 - \frac{1}{2} \left\| \mathcal{Q}_f^{-1/2}(\mathbf{p} - \mathbf{b}) \right\|_2^2 \right\}.$$

In this case, we can compute the maximizer in the above Hopf formula explicitly, which can be done numerically using the method of least squares.

Alternatively, this LQR problem can also be solved via the Riccati ODEs, which are given by

$$(3.7) \quad \dot{P}(t) = -P(t)^T B \mathcal{R}^{-1} B^T P(t), \quad \dot{\mathbf{q}}(t) = -P(t)^T B \mathcal{R}^{-1} (B^T \mathbf{q}(t) - \mathbf{a}),$$

$$\dot{r}(t) = -\frac{1}{2} \left\| \mathcal{R}^{-1/2}(B^T \mathbf{q}(t) - \mathbf{a}) \right\|_2^2$$

with initial conditions $P(0) = \mathcal{Q}_f$, $\mathbf{q}(0) = \mathbf{b}$, and $r(0) = 0$.

The above Hopf formula (3.6) is related to the learning problem (2.6) with a quadratic data fidelity term and quadratic regularization. Specifically, let $t = \lambda$ and $\mathbf{p} = \boldsymbol{\theta}$. Then, solving the above maximization problem (3.6) is equivalent to minimizing the following loss function with respect to $\boldsymbol{\theta} = [\theta_1, \dots, \theta_n]^T$:

$$(3.8) \quad \mathcal{L}(\boldsymbol{\theta}) = \frac{\lambda}{2} \left\| \mathcal{R}^{-1/2}(B^T \boldsymbol{\theta} - \mathbf{a}) \right\|_2^2 + \frac{1}{2} \left\| \mathcal{Q}_f^{-1/2}(\boldsymbol{\theta} - (\mathbf{b} + \mathcal{Q}_f \mathbf{x})) \right\|_2^2.$$

This loss function corresponds to a one-point linear regression problem, where the regularization term is given by $\frac{1}{2} \left\| \mathcal{Q}_f^{-1/2}(\boldsymbol{\theta} - (\mathbf{b} + \mathcal{Q}_f \mathbf{x})) \right\|_2^2$ and the data fitting term is given by $\frac{\lambda}{2} \left\| \mathcal{R}^{-1/2}(B^T \boldsymbol{\theta} - \mathbf{a}) \right\|_2^2$; i.e., set $N = 1$ in Figure 4. The minimizer $\boldsymbol{\theta}^*$ of \mathcal{L} is related to the solution of the Riccati equation via

$$(3.9) \quad \boldsymbol{\theta}^*(= \mathbf{p}^*) = \nabla_{\mathbf{x}} S(\mathbf{x}, \lambda) = P(\lambda) \mathbf{x} + \mathbf{q}(\lambda),$$

where \mathbf{p}^* is the minimizer in the Hopf formula (3.6). Note that if we only need to recover the minimizer $\boldsymbol{\theta}^*$, then we only need to solve two ODEs (namely, the ODEs for P, \mathbf{q}) since (3.9) does not depend on r .

3.3. Connection to multipoint regularized linear regression problems.

If the running cost is piecewise quadratic with $\mathcal{Q} = 0, \mathcal{N} = 0$ on each piece and the dynamics are piecewise linear with $A = 0$ on each piece, then we get a more general piecewise LQR problem in the form of (2.3) with piecewise running costs $L_i(\mathbf{u}) = \frac{1}{2} \mathbf{u}^T \mathcal{R}_i \mathbf{u} - \mathbf{a}_i^T \mathbf{u}$, where each $\mathcal{R}_i \in \mathbb{R}^{m \times m}$ is symmetric positive definite, piecewise

dynamics $f_i(\mathbf{u}) = B_i \mathbf{u}$, terminal cost $J(\mathbf{x}) = \frac{1}{2} \mathbf{x}^T \mathcal{Q}_f \mathbf{x} + \mathbf{b}^T \mathbf{x}$, and terminal time $t = \sum_{j=1}^N t_j$. Recall that this optimal control problem corresponds to the multitime HJ PDE (2.4) with Hamiltonians $H_i(\mathbf{p}) = \sup_{\mathbf{u} \in \mathbb{R}^m} \{ \langle -f_i(\mathbf{u}), \mathbf{p} \rangle - L_i(\mathbf{u}) \}$ and initial data J . Then, the solution to this LQR problem and corresponding HJ PDE is given by the following multitime Hopf formula:

$$(3.10) \quad S(\mathbf{x}, t_1, \dots, t_N) = \sup_{\mathbf{p} \in \mathbb{R}^n} \left\{ \langle \mathbf{x}, \mathbf{p} \rangle - \sum_{i=1}^N \frac{t_i}{2} \|\mathcal{R}_i^{-1/2} (B_i^T \mathbf{p} - \mathbf{a}_i)\|_2^2 - \frac{1}{2} \|\mathcal{Q}_f^{-1/2} (\mathbf{p} - \mathbf{b})\|_2^2 \right\}.$$

Define T_i to be

$$(3.11) \quad T_i = \sum_{j=1}^i t_j.$$

Then, the multitime HJ PDE and piecewise LQR problem can also be solved using the following Riccati equation: $S(\mathbf{x}, t_1, \dots, t_N) = \frac{1}{2} \mathbf{x}^T P(T_N) \mathbf{x} + \mathbf{q}(T_N)^T \mathbf{x} + r(T_N)$, where the function $P: [0, \infty) \rightarrow \mathbb{R}^{n \times n}$ takes values in the space of symmetric positive definite matrices and solves the following piecewise Riccati equation,

$$(3.12) \quad \begin{cases} \dot{P}(s) = -P(s)^T B_i \mathcal{R}_i^{-1} B_i^T P(s), & s \in (T_{i-1}, T_i), \\ P(0) = \mathcal{Q}_f, \end{cases}$$

the function $\mathbf{q}: [0, \infty) \rightarrow \mathbb{R}^n$ solves the following piecewise linear ODE,

$$(3.13) \quad \begin{cases} \dot{\mathbf{q}}(s) = -P(s)^T B_i \mathcal{R}_i^{-1} (B_i^T \mathbf{q}(s) - \mathbf{a}_i), & s \in (T_{i-1}, T_i), \\ \mathbf{q}(0) = -\mathbf{b}, \end{cases}$$

and the function $r: [0, \infty) \rightarrow \mathbb{R}$ solves the following piecewise ODE,

$$(3.14) \quad \begin{cases} \dot{r}(s) = -\frac{1}{2} \left\| \mathcal{R}_i^{-1/2} (B_i^T \mathbf{q}(s) - \mathbf{a}_i) \right\|_2^2, & s \in (T_{i-1}, T_i), \\ r(0) = 0. \end{cases}$$

The above multitime Hopf formula (3.10) can also be regarded as a linear regression problem with multiple data points, as illustrated in Figure 4. Let $\frac{1}{2} \|\mathcal{Q}_f^{-1/2} (\mathbf{p} - \mathbf{b})\|_2^2$ be the regularization term and $\frac{t_i}{2} \|\mathcal{R}_i^{-1/2} (B_i^T \mathbf{p} - \mathbf{a}_i)\|_2^2$ be the data fitting term at the i th data point. Then, the multitime LQR problem (3.10) is equivalent to the learning problem $\min_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta})$, where the loss function $\mathcal{L}: \mathbb{R}^n \rightarrow \mathbb{R}$ is given by

$$(3.15) \quad \mathcal{L}(\boldsymbol{\theta}) = \sum_{i=1}^N \frac{\lambda_i}{2} \left\| \mathcal{R}_i^{-1/2} (B_i^T \boldsymbol{\theta} - \mathbf{a}_i) \right\|_2^2 + \frac{1}{2} \left\| \mathcal{Q}_f^{-1/2} (\boldsymbol{\theta} - (\mathbf{b} + \mathcal{Q}_f \mathbf{x})) \right\|_2^2.$$

In section 4.1, we discuss a specific example of the learning problem (3.15), which is more readily recognizable as the standard linear regression problem. The minimizer $\boldsymbol{\theta}^*$ of the learning problem (3.15) (and \mathbf{p}^* of the Hopf formula (3.10)) is given by

$$(3.16) \quad \boldsymbol{\theta}^* (= \mathbf{p}^*) = \nabla_{\mathbf{x}} S(\mathbf{x}, t_1, \dots, t_N) = P(T_N) \mathbf{x} + \mathbf{q}(T_N).$$

For more details, we refer readers to [4]. Note that one could also use the Pontryagin maximum principle to compute the gradient of the solution to the HJ PDE, which gives the same result as in (3.16).

TABLE 1

Summary of the learning problems and corresponding numerical examples presented in this work. Each row summarizes which terms in the loss function (4.1), the Hopf formula, and the optimal control problem must be changed to match the context of various learning problems.

Terms in the loss function (4.1)	Terms In the Hopf formula	Terms in the optimal control problem	Learning problems
Φ_i, \mathbf{y}_i	Hamiltonian	Running cost, dynamics	Continual learning (section 5.1) Posttraining calibration (section 5.2)
λ_i	Time	Time	Posttraining calibration (section 5.2)
γ_k	Hamiltonian; Initial condition	Running cost, dynamics; Terminal cost	Hyperparameter tuning, Flow along the Pareto front (section 5.3)
θ_k^0	Spatial variable	Initial position	Generalized regularization functions (section 5.4)




FIG. 5. (See section 4). Mathematical formulation describing the connection between our model linear regression problem with quadratic regularization (top), the multitime Hopf formula (middle), and a piecewise LQR problem (bottom). Note that $T_i = \sum_{j=1}^i t_j$. The content of this illustration is a special case of the connection in Figure 4 using $\mathcal{R}_i = I$, $B_i = \Phi_i^T$, $\mathbf{a}_i = \mathbf{y}_i$, $\mathcal{Q}_f = \Gamma^{-1}$, $\mathbf{b} = 0$, and $\mathbf{x} = \Gamma \theta^0$. We use this connection to develop our Riccati-based methodology (section 4) and perform our numerical examples (section 5). The colors indicate the associated quantities between each problem. The solid-line arrows denote direct equivalences. The dotted arrows represent additional mathematical relations. Note: color appears only in the online article.

4. Methodology. In the previous section, we established a novel theoretical connection between multitime HJ PDEs, LQR problems, and regularized linear regression problems. In this section, we leverage this theoretical connection to design new algorithms for solving various learning problems. The connection between the learning problem of interest, the Hopf formula, and the corresponding optimal control problem is summarized in Table 1 and Figure 5.

4.1. Solving the regularized linear regression problem using Riccati ODEs. A linear regression problem with quadratic data fitting loss and quadratic regularization is formulated as follows. The goal of this learning problem is to fit N data points $(\mathbf{z}_i, \mathbf{y}_i) \in \mathbb{R}^M \times \mathbb{R}^m$ with the linear prediction model $\Phi_i \theta \approx \mathbf{y}_i$, where $\Phi_i = [\phi_1(\mathbf{z}_i), \dots, \phi_n(\mathbf{z}_i)] \in \mathbb{R}^{m \times n}$ is the matrix whose columns are the basis functions $\phi_j : \mathbb{R}^M \rightarrow \mathbb{R}^m$, $j = 1, \dots, n$, evaluated at \mathbf{z}_i and $\theta = [\theta_1, \dots, \theta_n]^T \in \mathbb{R}^n$ are unknown trainable coefficients. We learn $\theta \in \mathbb{R}^n$ by minimizing the following loss function

$$(4.1) \quad \mathcal{L}(\theta) = \frac{1}{2} \sum_{i=1}^N \lambda_i \|\Phi_i \theta - \mathbf{y}_i\|_2^2 + \frac{1}{2} \sum_{k=1}^n \gamma_k (\theta_k - \theta_k^0)^2,$$

where $\lambda_i \geq 0, i = 1, \dots, N$, are tunable hyperparameters for the data fitting losses, $\gamma_k > 0, k = 1, \dots, n$, are tunable hyperparameters for the regularization terms, and $\theta_k^0 \in \mathbb{R}$ is a prior on the unknown coefficients that biases θ_k to be close to θ_k^0 . Note that the above loss function (4.1) is strictly convex and, hence, it has a unique global minimizer. Since this loss function is quadratic in θ , it can be minimized exactly using the method of least squares. However, we explore a different approach for minimizing (4.1), which, in certain learning contexts, yields computational advantages over conventional approaches like the method of least squares.

Note that this learning problem is in the form of (3.15); i.e., set $\mathcal{R}_i = I_{m \times m}$, $B_i = \Phi_i^T$, $\mathbf{a}_i = \mathbf{y}_i$, $\mathcal{Q}_f = \Gamma^{-1}$, where $\Gamma \in \mathbb{R}^{n \times n}$ is the diagonal matrix whose i th entry is γ_i , $\mathbf{b} = 0$, and $\mathbf{x} = \Gamma \theta^0$. Then, this learning problem is related to an LQR problem, and we summarize this connection in Figure 5. Thus, this learning problem can alternatively be solved via the following Riccati ODEs:

$$(4.2) \quad \begin{cases} \dot{P}(t) = -P(t)^T \Phi_i^T \Phi_i P(t), & t \in (T_{i-1}, T_i), \\ \dot{\mathbf{q}}(t) = -P(t)^T \Phi_i^T (\Phi_i \mathbf{q}(t) - \mathbf{y}_i), & t \in (T_{i-1}, T_i) \end{cases}$$

with initial condition $P(0) = \Gamma^{-1}$, $\mathbf{q}(0) = 0$. Note that here we disregard the ODE for r since, for the learning problem, we are only concerned with the value of the minimizer θ^* , which only requires the values of P, \mathbf{q} (e.g., see (3.9), (3.16)), whereas recovering the minimal objective value $\mathcal{L}(\theta^*)$ (which is generally not needed in the context of learning) would require all three values P, \mathbf{q}, r . There are many numerical methods for solving the Riccati ODEs (4.2) in the literature. In this paper, we use the 4th order Runge-Kutta method (RK4) [6], but other methods could be used instead.

At first glance, solving this linear regression problem via the Riccati ODEs (4.2) may seem unnecessary given the existence of other well-established methods for minimizing (4.1) (e.g., the method of least squares). However, note that (4.2) is actually a sequence of Riccati ODEs. Thus, using our theoretical connection (and hence, these sequential Riccati ODEs to minimize (4.1)) means that we have the flexibility to handle sequential changes to the learning problem. In the remainder of section 4, we identify several applications in learning (summarized in Table 1), for which this Riccati-based approach yields computational advantages over traditional learning methods (especially when the number of data points N is large), and we describe how standard Riccati solvers can be adapted for these contexts. In Appendix D, we provide a more detailed comparison of the computational complexity of using RK4 and our Riccati-based methodology versus the method of least squares for these applications.

4.2. Adding or removing data. Since our theoretical connection gives us access to the Riccati ODEs (4.2), which are solved sequentially, we have the flexibility to handle sequential changes to the learning problem (4.1). In this section, we focus on the sequential addition or removal of data. Some related machine learning examples are provided in sections 5.1 and 5.2.

First, we discuss the addition of one data point; i.e., we increase the number of data points from N to $N + 1$. This case corresponds to updating our learned model as new data are collected, which is crucial for many practical machine learning applications. To add one data point, we adapt the Riccati ODEs (4.2) as follows. Adding the $(N + 1)$ th data point corresponds to adding the term $\frac{1}{2} \lambda_{N+1} \|\Phi_{N+1} \theta - \mathbf{y}_{N+1}\|_2^2$ in the loss function (4.1) or, equivalently, to adding the Hamiltonian $\frac{1}{2} \|\Phi_{N+1} \theta - \mathbf{y}_{N+1}\|_2^2$ to the multitime HJ PDE and the pieces $L_{N+1}(s, \mathbf{u}) = \frac{1}{2} \mathbf{u}^T \mathbf{u} - \mathbf{y}_{N+1}^T \mathbf{u}$ and $f(s, \mathbf{u}) =$

$\Phi_{N+1}^T \mathbf{u}, s \in (T_N, T_{N+1})$ to the running cost and dynamics, respectively, of the corresponding piecewise LQR problem. Thus, minimizing this new loss function is equivalent to solving the following Riccati ODE:

$$(4.3) \quad \begin{cases} \dot{\tilde{P}}(t) = -\tilde{P}(t)^T \Phi_j^T \Phi_j \tilde{P}(t), & t > 0, \\ \dot{\tilde{\mathbf{q}}}(t) = -\tilde{P}(t)^T \Phi_j^T (\Phi_j \tilde{\mathbf{q}}(t) - \mathbf{y}_j), & t > 0, \end{cases}$$

where the index $j = N+1$ and with initial condition $\tilde{P}(0) = P(T_N)$ and $\tilde{\mathbf{q}}(0) = \mathbf{q}(T_N)$, where $P(T_N)$ and $\mathbf{q}(T_N)$ are obtained from solving the learning problem (4.1) with N data points. Then, the solution to the new learning problem with an additional point is given by

$$(4.4) \quad \tilde{\boldsymbol{\theta}}^* = \tilde{P} \Gamma \boldsymbol{\theta}^0 + \tilde{\mathbf{q}},$$

where $\tilde{P} = \tilde{P}(\lambda_{N+1}) (= P(T_{N+1}))$ and $\tilde{\mathbf{q}} = \tilde{\mathbf{q}}(\lambda_{N+1}) (= \mathbf{q}(T_{N+1}))$ are the solution to (4.3), i.e., we have evolved the solution of the new corresponding multitime HJ PDE in the time variable t_{N+1} from $S(\mathbf{x}, t_1, \dots, t_N, 0)$ to $S(\mathbf{x}, t_1, \dots, t_N, \lambda_{N+1})$. Equivalently, we can also interpret adding one data point as solving a one-point linear regression problem (3.8) and its corresponding single-piece LQR problem (3.5).

Removing one data point (i.e., decreasing the number of data points from N to $N-1$) corresponds to calibrating our learned model by removing possible outliers and/or overly noisy data. To remove one data point, we reverse time and solve a terminal value Riccati ODE (C.1). The solution to the new learning problem is then given by (4.4), where \tilde{P} , $\tilde{\mathbf{q}}$ are given by the solution to the time-reversed Riccati ODEs. For more details and the mathematical derivation, see Appendix C.1.

Note that in both cases, the above approach only requires information about the data point to be added or removed and the results of the previous training. Thus, we can add and remove data without retraining on the entire dataset or requiring access to all of the previous data. In contrast, traditional approaches to minimizing (4.1) (e.g., the method of least squares) would require memory of all previous data and then retraining on the entire updated data set. Hence, our approach provides promising computational and memory savings over conventional methods. While, in the numerics, we use RK4 to solve the corresponding Riccati ODEs in both of these cases, we note that other numerical methods could also be applied to solve these Riccati ODEs and, hence, to achieve these computational and memory savings. For example, in Appendix E, we discuss how recursive least squares can be applied to these cases.

4.3. Hyperparameter tuning. In the loss function (4.1), we have three types of hyperparameters: the weights λ_i of the data fitting terms, the weights γ_k of the regularization term, and the bias $\boldsymbol{\theta}^0$ on the trainable coefficients $\boldsymbol{\theta}$. In this section, we discuss how we adapt the Riccati ODEs to tune each of these hyperparameters.

First, we discuss tuning the weights of the data fitting terms, e.g., as in posttraining calibration and federated learning (section 5.2). Consider changing the weight on the i th data fitting term from λ_i to $\tilde{\lambda}_i$, which corresponds to changing the time $t_i = \lambda_i$ to $t_i = \tilde{\lambda}_i$ in the multitime HJ PDE and corresponding piecewise LQR problem. Then, the methodology is similar to that in section 4.1.

If $\tilde{\lambda}_i > \lambda_i$, then the solution to the learning problem with the new data fitting weight $\tilde{\lambda}_i$ is given by (4.4), where $\tilde{P} = \tilde{P}(\tilde{\lambda}_i - \lambda_i)$ and $\tilde{\mathbf{q}} = \tilde{\mathbf{q}}(\tilde{\lambda}_i - \lambda_i)$ are the solutions to the Riccati ODEs (4.3) with $j = i$, at time $(\tilde{\lambda}_i - \lambda_i)$, and with initial condition $\tilde{P}(0) = P(T_N)$ and $\tilde{\mathbf{q}}(0) = \mathbf{q}(T_N)$, where $P(T_N)$, $\mathbf{q}(T_N)$ are obtained from solving the original learning problem (4.1) with the original value of λ_i . Similarly, if $\lambda_i > \tilde{\lambda}_i$, then

the solution to the learning problem with the new data fitting weight $\tilde{\lambda}_i$ is given by (4.4), where $\tilde{P} = \tilde{P}(0)$ and $\tilde{\mathbf{q}} = \tilde{\mathbf{q}}(0)$ are the solutions to the time-reversed Riccati ODEs (C.1) with $j = i$, at time 0, and with terminal condition $\tilde{P}(\lambda_i - \tilde{\lambda}_i) = P(T_N)$ and $\tilde{\mathbf{q}}(\lambda_i - \tilde{\lambda}_i) = \mathbf{q}(T_N)$, where $P(T_N), \mathbf{q}(T_N)$ are obtained from solving the original learning problem (4.1) with the original value of λ_i .

Next, we discuss tuning the weights of the regularization terms, e.g., as in the hyperparameter tuning example in section 5.3. Note that in the original formulation of the Riccati ODEs (4.2), the regularization weights appear in the initial condition for P . Changing a hyperparameter γ_k by changing this initial condition would require resolving the entire sequence of Riccati ODEs and hence retraining on the entire dataset. However, since we have freedom in how we formulate the corresponding multitime HJ PDE, we can avoid this retraining as follows. Note that here, instead of sequentially changing each regularization weight γ_k one-by-one, we can actually update all of the regularization weights at once.

We update the regularization weights in two steps. First, we update the weights whose values we increase by adding an additional Hamiltonian to the multitime HJ PDE. This yields an initial value Riccati ODE (C.2). Second, we update the remaining weights by reinterpreting the problem as a terminal condition, single-time HJ PDE. We then evolve this HJ PDE backward in time using a terminal condition Riccati ODE (C.3). Finally, the minimizer to the new learning problem resulting from changing all of the regularization weights is given by the solutions to the time-reversed Riccati ODE (C.3) at time 0. For more details, see Appendix C.2.

Finally, we discuss tuning the bias $\boldsymbol{\theta}^0$ on the trainable coefficients $\boldsymbol{\theta}$, e.g., as in the generalized regularization example in section 5.4. In the context of the learning problem, we interpret $\boldsymbol{\theta}^0$ as the center of a Gaussian prior on $\boldsymbol{\theta}$. Using our theoretical connection, changing $\boldsymbol{\theta}^0$ is also equivalent to evaluating the corresponding multitime HJ PDE and piecewise LQR problem at a different point in space. Specifically, changing $\boldsymbol{\theta}^0$ to $\tilde{\boldsymbol{\theta}}^0$ means that instead of evaluating the multitime HJ PDE and piecewise LQR problem at $\mathbf{x} = \Gamma\boldsymbol{\theta}^0$, we evaluate them at $\mathbf{x} = \Gamma\tilde{\boldsymbol{\theta}}^0$. Then, the solution to the learning problem with the new bias $\tilde{\boldsymbol{\theta}}^0$ is given by

$$(4.5) \quad P(T_N)\Gamma\tilde{\boldsymbol{\theta}}^0 + \mathbf{q}(T_N),$$

where $P(T_N)$ and $\mathbf{q}(T_N)$ are obtained from solving the original learning problem (4.1) with the original value of $\boldsymbol{\theta}^0$. In other words, shifting the bias involves neither retraining nor access to any of the data points nor solving Riccati ODEs. Instead, we only require some matrix multiplication and addition involving the results of the previous training and the new bias $\tilde{\boldsymbol{\theta}}^0$.

4.4. General convex regularization functions. So far, we have only considered quadratic regularizations. In this section, we will consider the linear regression problem with an arbitrary convex regularization term. Specifically, consider the general regularization term $R(\boldsymbol{\theta}) - \langle \mathbf{x}, \boldsymbol{\theta} \rangle$, where $R: \mathbb{R}^n \rightarrow \mathbb{R}$ is a convex function. Then, the loss function of the learning problem is given by

$$(4.6) \quad \mathcal{L}(\boldsymbol{\theta}) = \frac{1}{2} \sum_{i=1}^N \lambda_i \|\Phi_i \boldsymbol{\theta} - \mathbf{y}_i\|_2^2 + R(\boldsymbol{\theta}) - \langle \mathbf{x}, \boldsymbol{\theta} \rangle.$$

This learning problem corresponds to a multitime HJ PDE (2.4), where the i th Hamiltonian is $H_i(\mathbf{p}) = \frac{1}{2} \lambda_i \|\Phi_i \mathbf{p} - \mathbf{y}_i\|_2^2$ and the initial condition is $J = R^*$, which is the

Fenchel–Legendre transform of R . The corresponding optimal control problem is given by (2.3), with terminal time $\sum_{i=1}^N \lambda_i$, piecewise running costs $L_i(\mathbf{u}) = \frac{1}{2} \|\mathbf{u}\|^2 - \langle \mathbf{y}_i, \mathbf{u} \rangle$, and piecewise dynamics $f_i(\mathbf{u}) = \Phi_i^T \mathbf{u}$ on $(\sum_{j=1}^{i-1} \lambda_j, \sum_{j=1}^i \lambda_j]$, and terminal cost R^* . Then, the solution $\boldsymbol{\theta}^*$ to the learning problem (4.6) is equivalent to the spatial gradient $\nabla_{\mathbf{x}} S(\mathbf{x}, \lambda_1, \dots, \lambda_N)$ of the solution to the multitime HJ PDE, and the minimal value of the loss function (4.6) is equivalent to $-S(\mathbf{x}, \lambda_1, \dots, \lambda_N)$.

Since the loss function (4.6) is convex, it can be minimized using any appropriate convex optimization algorithm. In this paper, we demonstrate how our Riccati-based approach can be combined with PDHG [7] to solve this learning problem. Note that PDHG can be applied to this learning problem (4.6) as long as the proximal point of R or R^* is numerically computable. For instance, R could be a quadratic function $\langle \cdot, \mathcal{M} \cdot \rangle$, a quadratic norm $\sqrt{\langle \cdot, \mathcal{M} \cdot \rangle}$, $\|\cdot\|_1$, $\|\cdot\|_1^2$, $\|\cdot\|_\infty^2$, a polynomial function, or the Fenchel–Legendre transformation of any of these functions [17]. To compute the solution $\boldsymbol{\theta}^*$ of this learning problem, PDHG iterates the following,

$$(4.7) \quad \begin{cases} \boldsymbol{\theta}^{\ell+1} = \operatorname{argmin}_{\boldsymbol{\theta} \in \mathbb{R}^n} \frac{1}{2} \sum_{i=1}^N \lambda_i \|\Phi_i \boldsymbol{\theta} - \mathbf{y}_i\|_2^2 + \frac{1}{2\sigma_\theta} \left\| \boldsymbol{\theta} - (\boldsymbol{\theta}^\ell - \sigma_\theta(\mathbf{w}^\ell - \mathbf{x})) \right\|_2^2, \\ \bar{\boldsymbol{\theta}}^{\ell+1} = 2\boldsymbol{\theta}^{\ell+1} - \boldsymbol{\theta}^\ell, \\ \mathbf{w}^{\ell+1} = \operatorname{argmin}_{\mathbf{w} \in \mathbb{R}^n} R^*(\mathbf{w}) + \frac{1}{2\sigma_w} \left\| \mathbf{w} - (\mathbf{w}^\ell + \sigma_w \bar{\boldsymbol{\theta}}^{\ell+1}) \right\|_2^2, \end{cases}$$

where σ_θ, σ_w are positive step-size parameters satisfying $\sigma_\theta \sigma_w < 1$. In each iteration, we need to solve two minimization problems. Updating $\boldsymbol{\theta}^\ell$ is equivalent to solving the learning problem with quadratic regularization. Specifically, updating $\boldsymbol{\theta}^\ell$ is the same as changing the bias $\boldsymbol{\theta}^0$ in the loss function (4.1). Therefore, we only need to solve the Riccati ODEs (4.2) once (to compute $\boldsymbol{\theta}^1$) and every other iterate can be computed using (4.5). Updating \mathbf{w}^ℓ is equivalent to computing the proximal point of R^* . Depending on R^* , there may already exist efficient solvers or explicit formulas for computing its proximal point.

Note that with nonquadratic regularization, if we change the weights λ_k of the data fitting losses, change \mathbf{x} (which is related to the bias on $\boldsymbol{\theta}$), or add or remove data points, then we will need to restart PDHG to solve the resulting learning problem with these new hyperparameter values or updated datasets. However, we can reuse some of the computations between runs. Namely, we only need to solve the full sequence of Riccati ODEs (4.2) once (to compute $\boldsymbol{\theta}^1$ in the first run of PDHG). Then, the solution to those Riccati ODEs can be reused in combination with the methods described in sections 4.2 and 4.3 (according to how the learning problem is modified) in subsequent iterations and runs of PDHG.

5. Numerical examples. In this section, we apply the Riccati-based methodology presented in section 4 to four test problems from machine learning to demonstrate the versatility and potential computational advantages of our new approach. In each example, we use RK4 with double precision to solve the Riccati ODEs when applying our Riccati-based methodology. The connections between the examples presented in this section, the Hopf formula, and the corresponding optimal control problems can be found in Table 1. We note that in this work we use two metrics to evaluate our results quantitatively: the ℓ_1 -norm (defined as $\|x\|_1 = \sum_{i=1}^n |x_i|$ for $x \in \mathbb{R}^n$) for the finite-dimensional minimizer and the L^2 -norm (defined as $\|f\|_2 = \left(\int_{x \in \mathbb{R}^n} |f(x)|^2 dx \right)^{1/2}$ for $f: \mathbb{R}^n \rightarrow \mathbb{R}$) for the inferred functions. The L^2 -norm for functions is approximated using the trapezoidal rule with a uniform grid. Supplementary details of the numerical experiments can be found in Appendix G. Code for all examples will be made publicly available at <https://github.com/ZongrenZou/HJPDE4SciML>.

5.1. Function approximation in continual learning. In this section, we test our Riccati-based approach (as described in section 5.1) on a pedagogical function approximation example in continual learning [39, 28, 47] to demonstrate its computational and memory advantages. Under the continual learning framework, data are accessed in a stream and the trainable model parameters are updated incrementally as new data become available. In some cases, the historical data may also become inaccessible after new data are received, which can often lead to catastrophic forgetting [28, 39], which refers to the abrupt degradation in performance of learned models on previous tasks upon training on new tasks. In this example, we show how our Riccati-based approach naturally coincides with the continual learning framework, while also inherently avoiding catastrophic forgetting even if the historical data are inaccessible.

Our goal is to regress the function $y(x) = \sin(10x)$ given noisy data $\{x_i, y_i \approx y(x_i)\}_{i=1}^N$, where the data are corrupted by additive Gaussian noise with relatively large noise scale. Hence, a large amount of data is required in order to obtain an accurate inference. In this example, we follow the continual learning framework and assume that the data are accessed in a stream, so that each new data point must be incorporated into the learned model as soon as it becomes available. Specifically, each new data point comes from a uniformly random sample point $x \in [0, 10]$ and is corrupted by additive Gaussian noise with noise scale one. We regress $y(x)$ using the linear model $y(x) = \sum_{k=1}^n \theta_k \phi_k(x)$, where $n = 10$ and

$$\{\phi_k(x)\}_{k=1}^n = \{1, x, x^2, x^3, \sin(\tau), \sin(5x), \sin(8x), \sin(9x), \sin(10x), \sin(12x)\}.$$

We learn the coefficients $\boldsymbol{\theta} = [\theta_1, \dots, \theta_n]^T$ by minimizing the following loss function,

$$(5.1) \quad \mathcal{L}(\boldsymbol{\theta}) = \frac{1}{2} \sum_{i=1}^N \lambda_i \left| \sum_{k=1}^n \theta_k \phi_k(x_i) - y_i \right|^2 + \frac{1}{2} \sum_{k=1}^n \gamma_k |\theta_k|^2,$$

where $\lambda_i, i = 1, \dots, N$ and $\gamma_k, k = 1, \dots, n$ are weights for the data loss and ℓ_2 regularization terms, respectively, and we update our learned coefficients every time a new data point is available. In our numerical experiments, we set $\lambda_i = 1 \forall i$ and $\gamma_k = 100 \forall k$.


Although this loss function (5.1) could be minimized using conventional machine learning techniques (e.g., the method of least squares), these methods typically require access to and training on the entire dataset $\{(x_i, y_i)\}_{i=1}^N$, which conflicts with the assumptions in continual learning. However, note that this loss function (5.1) is of the form (4.1). Thus, we can instead apply our Riccati-based approach (as described in sections 4.1 and 4.2) to solve this learning problem. In particular, the methodology described in section 4.2 matches the data streaming paradigm of continual learning; we incrementally update our learned coefficients by considering each new data point as the time evolution of a corresponding multitime HJ PDE. As a result, in contrast to conventional machine learning methods, our Riccati-based approach does not require storage of previous data points, and its memory and computational complexity is constant for each added data point. As such, our Riccati-based methodology may be well suited to online learning applications. Additionally, since this time evolution of the HJ PDE (i.e., the addition of a new data point) requires knowledge about the solution to the HJ PDE at the previous time (i.e., the results of the previous training), our Riccati-based approach also inherently avoids catastrophic forgetting.

The results of applying our Riccati-based method to solve this learning problem (5.3) are shown in Table 2 and Figure 6. Figure 6(a) depicts the evolution of the coefficients $\{\theta_k\}_{k=1}^n$ as more data points become available. We observe that the learned


TABLE 2

Errors in computing the minimizer θ^* of the function approximation loss (5.1) using our Riccati-based approach. We use RK4 to solve the Riccati ODEs (4.2) and (4.3) with double precision and various step sizes h . The reference is obtained by using the method of least squares to minimize the loss function (5.1) directly.

	$h = 0.001$	$h = 0.0005$	$h = 0.0001$
ℓ_1 error of θ^*	2.7814×10^{-7}	4.9461×10^{-10}	3.0801×10^{-12}
ℓ_1 relative error of θ^*	2.7196×10^{-7}	4.8361×10^{-10}	3.0116×10^{-12}



(a)



(b)

FIG. 6. Evolution and continual learning of the function approximation learned using our Riccati-based approach as more data become available. (a) shows the evolution of the learned coefficients $\theta_k, k = 1, \dots, n$, and the relative L^2 error of the inferred y as more data are incorporated into the model; the horizontal dotted lines denote the exact reference values. (b) shows the inferences of y after the 200th, 1,000th, and 50,000th noisy data points become available. Our Riccati-based approach allows us to incrementally update the learned coefficients as more data become available without accessing the previous data or retraining on the entire dataset, which provides both memory and computational advantages over conventional learning methods. Note: color appears only in the online article.

coefficients $\theta_k, k = 1, \dots, n$, converge to their true values as more data points are incorporated into the model. Figure 6(b) displays the inferences obtained when the 200th, 1,000th, and 50,000th noisy data points become available, which demonstrates that our Riccati-based approach is capable of real-time inferences without catastrophic forgetting, even though each inference is made using only one data point. As shown, the regression using data with high-level noise becomes more accurate as more data are incorporated into the learning process.

Table 2 displays the numerical errors of the minimizer θ^* of the loss function (5.1) obtained using the Riccati-based methodology from section 4 after the last data point becomes available. We observe that the accuracy of our approach increases as we decrease the step size h of RK4, which indicates that the errors of θ^* stem from the

accuracy of RK4 in solving the corresponding Riccati ODEs. The reference solution is obtained by minimizing (5.1) directly using the method of least squares and assuming that all $N = 50,000$ data points are accessible.

5.2. 1D steady-state reaction-diffusion equation and posttraining calibration. In this example, we use our Riccati-based approach to apply posttraining calibrations when solving a PDE. Specifically, we leverage the methodology in section 4.2 to add or remove data and the methodology in section 4.3 to enforce the boundary conditions of the PDE without retraining the entire learned model. Consider the following 1D steady-state reaction-diffusion equation,

$$(5.2) \quad \begin{cases} D \frac{\partial^2 u}{\partial x^2}(x) + \kappa u(x) = f(x), x \in [0, 1], \\ u(0) = u(1) = 0, \end{cases}$$

where $D = 0.01$ is the diffusion coefficient, $\kappa = -1$, and $f(x)$ is the source term of which noisy measurements $\{(x_i, f_i \approx f(x_i))\}_{i=1}^N \subset \mathbb{R} \times \mathbb{R}$ are available. We consider the scenario where regular training has been employed but with either insufficient data or sufficient data with outliers, both of which yield inaccurate inferences of the solution. We further assume that extra information is provided after the regular training and seek to perform posttraining calibrations to incorporate this extra information into the already-trained models without losing information from the original training. In the literature, it is well established that posttraining calibrations can significantly increase the performance of deployed machine learning methods [40, 53]. However, designing computationally efficient methods for performing these postcalibrations is still of great interest.

In this example, we solve this PDE (5.2) by reformulating the PDE as an optimization problem [41, 44, 22]. We use a linear model to approximate the solution, i.e., $u(x) = \sum_{k=1}^n \theta_k \phi_k(x)$, where $n = 21$ and $\{\phi_k(x)\}_{k=1}^n = \{1\} \cup \{\sin(2l\pi x), \cos(2l\pi x)\}_{l=1}^{(n-1)/2}$ are the truncated Fourier basis functions on $[0, 1]$. We learn the coefficients $\theta = [\theta_1, \dots, \theta_n]^T$ of the linear model by minimizing the following loss,

$$(5.3) \quad \begin{aligned} \mathcal{L}(\theta) = & \frac{1}{2} \sum_{i=1}^N \lambda_i \left| D \sum_{k=1}^n \theta_k \frac{\partial^2 \phi_k}{\partial x^2}(x_i) + \kappa \sum_{k=1}^n \theta_k \phi_k(x_i) - f_i \right|^2 \\ & + \frac{1}{2} \lambda_b \left| \sum_{k=1}^n \theta_k \phi_k(0) - 0 \right|^2 + \frac{1}{2} \lambda_b \left| \sum_{k=1}^n \theta_k \phi_k(1) - 0 \right|^2 + \frac{1}{2} \sum_{k=1}^n \gamma_k |\theta_k|^2, \end{aligned}$$

where $\lambda_i, i = 1, \dots, N$, λ_b , and $\gamma_k, k = 1, \dots, n$, are balancing weights for the PDE residual, the boundary conditions, and the regularization term, respectively. In our numerical experiments, we assume the exact solution to be $u(x) = \sin^3(2\pi x)$ (and f to be defined by (5.2), accordingly) and the noise to be additive Gaussian with zero mean and standard deviation 0.1. For the regular training, we set $\lambda_i = 1$, $\lambda_b = 1$, and $\gamma_k = 1$ and apply our Riccati-based approach (see section 4.1) to get our original estimate of the minimizer θ^* of the loss function (5.3). Using our Riccati-based approach yields an ℓ_1 error of 8.9950×10^{-10} and relative ℓ_1 error of 1.4288×10^{-9} in θ^* , where the reference is obtained by minimizing (5.3) directly using the method of least squares.

In the leftmost column of Figure 7, we see that the accuracy of both u and f as inferred by the regular training is impaired by a lack of data around the highest peak and lowest valley of the exact functions. To compensate, we first calibrate our model by adding some new noisy measurements of f in these regions where the data points are sparse. This calibration uses the methodology described in section 4.2, and the

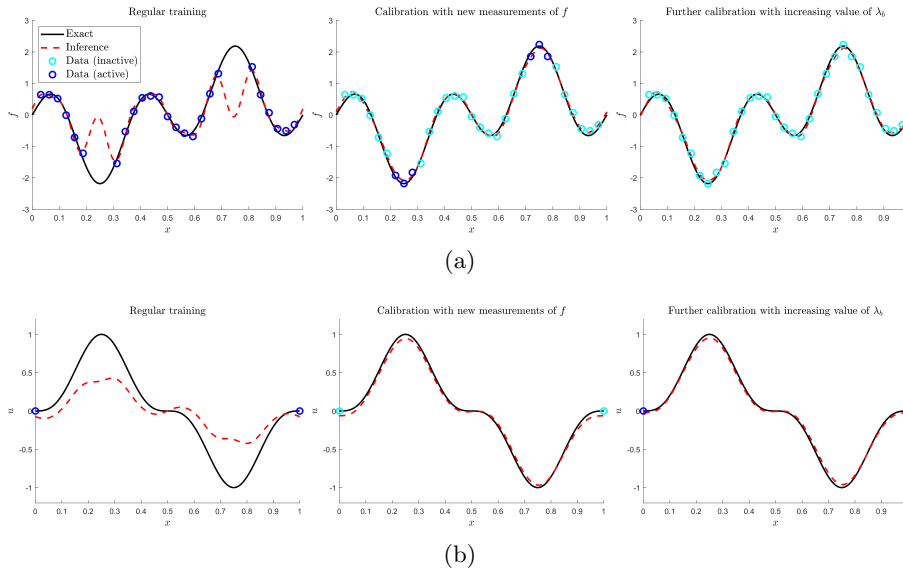


FIG. 7. Results of solving the 1D steady-state reaction-diffusion equation (5.2) with noisy measurements of the source term f in the domain and noiseless measurements of the solution u on the boundary. (a) results for f ; (b) results for u . Left: results of regular training; middle: calibrating the results with some additional noisy measurements of f ; right: calibrating the results further by enforcing the boundary conditions by increasing the value of λ_b in the loss function (5.3). The regular training uses our Riccati-based method in section 4.1 to minimize (5.3), while the calibrations use the adaptations of our method in section 4.2. Calibrations are employed without retraining or access to the data from the previous training, which demonstrates the advantages in both memory storage and computational complexity of our Riccati-based approach over conventional machine learning methods. Note: color appears only in the online article.

TABLE 3

Errors in solving the 1D steady-state reaction-diffusion equation (5.2) and regressing f at different stages of training, using our Riccati-based approach. Regular training is done with insufficient data and provides inaccurate inferences. The first calibration adds new measurements of f , and the second calibration increases the weights λ_b of the boundary conditions in (5.3). Both calibration steps successfully improve the accuracy of our inferences without requiring retraining on or access to the previous data. Qualitative results can be found in Figure 7.

	Regular training	First calibration	Second calibration
Relative L^2 error of f	66.40%	7.15%	6.49%
Relative L^2 error of u	57.29%	7.35%	5.57%

results are shown in the middle column of Figure 7. Next, we note that the inferred u still disagrees with the exact solution at the boundary points. Hence, we further calibrate our model by increasing the value of the boundary weight λ_b from 1 to 10 to enforce the boundary conditions of the PDE. This calibration is done using the methodology described in section 4.3 and the results of this second calibration are presented in the rightmost column of Figure 7. In both cases, we observe that the calibrations successfully improve the accuracy of the learned model. These results are also reflected in the relative L^2 errors shown in Table 3.

Note that our Riccati-based approach allows us to perform each of these calibration steps using only the new or changed values in that step and the results of the previous training step. In other words, each step of this training process (including the

original training and each subsequent calibration step) is done without sharing data between training steps, which exactly matches the framework of federated learning [30]. Thus, our methodology may be relevant to distributed training or collaborative learning applications, where data privacy is of concern.

Next, we discuss another posttraining calibration technique. In this case, we assume the data for the regular training are sufficient but contain outliers due to large noise. We again assume the regular training is performed using the Riccati-based approach from section 4.1. Then, we eliminate these outliers using the methodology described in section 4.2, which only requires knowledge about the outliers to be removed and the results of the regular training. The results of this posttraining outlier removal show that eliminating these points successfully improves the accuracy of the learned models (see Appendix F, Figure 9). Finally, in Appendix F, Figure 10, we also provide a large-scale continual learning example for solving (5.2) to further demonstrate the performance of our Riccati-based approach in large data settings.

5.3. Poisson equation using PINNs and transfer learning. In this example, we demonstrate the versatility of our Riccati-based method by combining it with existing machine learning techniques to fit the last layer of a PINN. We also show that when we perform hyperparameter tuning by solving the associated Riccati ODEs (section 4.3), we not only provide the solution to the updated problem but also a continuum of solutions along a 1D curve on the Pareto front of the data fitting losses and regularization. Consider the two-dimensional (2D) Poisson equation with Dirichlet boundary conditions, given by

$$(5.4) \quad \begin{cases} \frac{\partial^2 u}{\partial x^2}(x, y) + \frac{\partial^2 u}{\partial y^2}(x, y) = f(x, y), & (x, y) \in \Omega, \\ u(x, y) = 0, & (x, y) \in \partial\Omega, \end{cases}$$

where $\Omega := [0, 1]^2$ and f is a source term. We solve this equation using transfer learning and PINNs. Consider the scenario where we only have access to measurements $\{(x_i, y_i, f_i = f(x_i, y_i))\}_{i=1}^N$ of f at limited sampling points. Transfer learning compensates for this lack of knowledge by transferring the knowledge from models pretrained on similar problems to solve this new problem of interest. Here, we learn the linear model $u(x, y) = \sum_{k=1}^n \theta_k \phi_k(x, y)$, where each basis function $\phi_k(x, y), k = 1, \dots, n$ is the PINN solution of a neural network pretrained to solve the 2D Poisson equation (5.4) with similar source terms. Transfer learning using pretrained neural networks as basis functions, as we do here, has recently grown in popularity in the scientific machine learning community [18, 52, 21] and has been shown to provide efficient yet accurate inferences even given very limited data [52]. We learn the coefficients $\boldsymbol{\theta}$ of our linear model by minimizing the following PINN-type loss,

$$(5.5) \quad \mathcal{L}(\boldsymbol{\theta}) = \frac{1}{2} \sum_{i=1}^N \lambda_i \left| \sum_{k=1}^n \theta_k \left(\frac{\partial^2 \phi_k}{\partial x^2} + \frac{\partial^2 \phi_k}{\partial y^2} \right) (x_i, y_i) - f_i \right|^2 + \frac{1}{2} \sum_{k=1}^n \gamma_k |\theta_k|^2,$$

where $\lambda_i = 1, i = 1, \dots, N$, and $\gamma_k = \gamma, k = 1, \dots, n$, are weights for the data fitting and ℓ_2 -regularization terms, respectively. Note that minimizing (5.5) with respect to $\boldsymbol{\theta}$ is equivalent to fitting the last layer of a neural network given previous, pretrained nonlinear layers as the basis functions.

In our numerical experiments, we use transfer learning to solve the 2D Poisson equation (5.4) with source term $f(x, y) = \sin(2.5\pi x)\sin(2.5\pi y)$ using $n = 100$ basis functions and $N = 100$ random measurements of f . We use the multihead PINN method [52] to obtain the basis functions, which correspond to the shared nonlinear

TABLE 4

Errors of the minimizer θ^* of (5.5) and the solution u to the 2D Poisson equation (5.4) using transfer learning and our Riccati-based approach. The reference for θ^* is given by minimizing (5.5) with the corresponding value of γ directly using the method of least squares, and the reference for u is given by solving (5.4) using a finite difference method. Since we decrease γ incrementally from 1 to 1e-5, the error of θ^* accumulates due to successive applications of RK4.

	$\gamma = 10^0$	$\gamma = 10^{-1}$	$\gamma = 10^{-2}$	$\gamma = 10^{-3}$	$\gamma = 10^{-4}$	$\gamma = 10^{-5}$
ℓ_1 error of θ^*	5.9707×10^{-10}	4.2905×10^{-8}	1.0725×10^{-6}	1.7235×10^{-5}	2.0490×10^{-4}	5.8749×10^{-3}
L^2 relative error of u	6.5247%	5.6793%	3.0390%	1.7280%	1.1662%	1.4736%

hidden layers of the pretrained PINN solutions to (5.4) with source terms $f(x, y) = \sin(k\pi x) \sin(k\pi y)$, $k = 1, 2, 3, 4$. We note that the boundary condition is hard-encoded in the basis functions and, hence, in (4.1), no penalty term involving the boundary condition is included. We then solve the learning problem (5.5) using our Riccati-based method from section 4.1.

In Table 4, we compare the errors of the minimizer θ^* of (5.5) and the solution u of the PDE (5.4) as we decrease the weight $\gamma (= \gamma_k \forall k)$ of the regularization term from 1 to 1e-5. The reference for θ^* is obtained from minimizing (5.5) directly for each value of γ using the method of least squares. The reference for u is computed using a finite difference method with a five-point stencil and a 257×257 uniform grid on Ω to solve (5.4). The same grid is used to evaluate our trained models. We originally minimize (5.5) using $\gamma = 1$ and the Riccati-based approach in section 4.1. We then compute the solutions for the other values of γ by incrementally decreasing γ by a factor of 10 and using the methodology from section 4.3 to reuse the results of training with the previous value of γ to compute the solution for the new value of γ . Consequently, we see that the error in θ^* increases as we decrease γ due to error accumulation from repeated applications of RK4. However, the error of u generally decreases as we decrease γ with the lowest error being achieved when $\gamma = 1e-4$.

From the results in Table 4, we see that our choice of the hyperparameter γ can greatly influence the accuracy of our learned model. Note that since we fix $\lambda_i = 1 \forall i$ and $\gamma_k = \gamma \forall k$, we can view (5.5) as a biobjective loss, where the two objectives are the weighted data fitting term $\frac{1}{2} \sum_{i=1}^N \left| \sum_{k=1}^n \theta_k \left(\frac{\partial^2 \phi_k}{\partial x^2} + \frac{\partial^2 \phi_k}{\partial y^2} \right) (x_i, y_i) - f_i \right|^2$ and the regularization term $\frac{1}{2} \sum_{k=1}^n |\theta_k|^2$. To better understand the effects of tuning γ , in Figure 8, we explore the Pareto front of these two objectives. Traditional scalarization-based approaches for computing the Pareto front typically rely on discrete samplings of the Pareto front corresponding to discrete choices of γ [25]. While our Riccati-based methodology from section 4.3 also recovers discrete points on the Pareto front corresponding to particular choices of γ , note that when we change γ , e.g., from $\gamma = 1$ to $\gamma = 0.1$, we also recover a 1D curve along the Pareto front corresponding to every $\gamma \in [0.1, 1]$. We obtain this 1D curve theoretically via the flow of solutions obtained from the corresponding Riccati ODEs and numerically via the intermediate steps of RK4. The left plot of Figure 8 shows the 1D curve along the Pareto front recovered by our Riccati-based approach (although note that in this example, the Pareto front is also 1D and hence is equivalent to the exposed 1D curve), where the flow of solutions corresponding to decreasing γ is represented by the arrows. Thus, although in general our methodology cannot compute the entire Pareto front, every time we change the value of the hyperparameters, our approach recovers a continuous 1D curve along the




FIG. 8. Results of changing the regularization weight γ when solving the 2D Poisson equation using PINNs and transfer learning. We incrementally decrease the value of γ by evolving the corresponding Riccati ODEs backwards in time. As a result, we obtain a flow of solutions, the direction of which is represented by the arrows in each figure. In the left figure, this flow of solutions gives us that each change in the value of γ from $\hat{\gamma}$ to $\tilde{\gamma}$ results in the recovery of every point of the Pareto front along the 1D curve parameterized by $\gamma \in [\hat{\gamma}, \tilde{\gamma}]$. Note: color appears only in the online article.

Pareto front. In the right plot of Figure 8, we also visualize how the L^2 error of our learned solution u changes as we decrease γ .

5.4. Identifying the dynamics of the Kraichnan–Orszag system from data. In this example, we demonstrate the versatility of our Riccati-based approach by showing how it can be combined with existing methods to solve more general problems (see section 4.4). Consider the Kraichnan–Orszag–(K–O) system [48, 53, 51]

$$(5.6) \quad \begin{cases} \frac{dx_1}{d\tau} = x_2 x_3, \\ \frac{dx_2}{d\tau} = x_1 x_3, \\ \frac{dx_3}{d\tau} = -2x_1 x_2 \end{cases}$$

with initial conditions $x_1(0) = 1, x_2(0) = 0.8, x_3(0) = 0.5$. Our goal is to identify the dynamics (the right-hand side of (5.6)) of the K–O system using measurements of x_i and $\frac{dx_i}{d\tau}$, $i = 1, 2, 3$, at different times. We identify the dynamics by learning the linear models $\frac{dx_i}{d\tau} = \sum_{k=1}^n \theta_k^i \phi_k$, $i = 1, 2, 3$. Following the general framework of the SINDy method [5], we use the following quadratic basis functions ($n = 10$) for the dynamics:

$$\{\phi_k(x_1, x_2, x_3)\}_{k=1}^n = \{1, x_1, x_2, x_3, x_1^2, x_2^2, x_3^2, x_1 x_2, x_2 x_3, x_1 x_3\},$$

and impose ℓ_1 -regularization on θ to promote sparse identification of the dynamics. Then, we learn each θ^i independently and in parallel by minimizing the loss functions

$$(5.7) \quad \mathcal{L}_i(\theta^i) = \frac{1}{2} \sum_{j=1}^N \lambda_j \left[\left(\frac{dx_i}{d\tau} \right)_j - \sum_{k=1}^n \theta_k^i \phi_k((x_1)_j, (x_2)_j, (x_3)_j) \right]^2 + \sum_{k=1}^n \gamma_k |\theta_k^i|,$$

where \mathcal{L}_i denotes the loss function for equation i , $(x_i)_j$ and $\left(\frac{dx_i}{dt}\right)_j$ denote the measurements of x_i and $\frac{dx_i}{d\tau}$, respectively, at time τ_j , $i = 1, 2, 3, j = 1, \dots, N$. Note that (5.7) corresponds to setting $R = \|\cdot\|_1$ and $\mathbf{x} = 0$ in the linear regression problem (4.6). Thus, solving this learning problem is equivalent to evaluating the solution

TABLE 5

Results of sparse identification of the K-O system (5.6) using PDHG to minimize the ℓ_1 -regularized losses (5.7). The true solution is 0 for all entries, except $\theta^{1,*} = 1$ for x_2x_3 , $\theta^{2,*} = 1$ for x_1x_3 , and $\theta^{3,*} = -2$ for x_1x_2 . We recover the dynamics reasonably well, albeit with some slight misidentification of $\theta^{2,*}, \theta^{3,*}$.

	1	x_1	x_2	x_3	x_1^2	x_2^2	x_3^2	x_1x_2	x_2x_3	x_1x_3
$\theta^{1,*}$	0	0	0	0	0	0	0	0	0.9931	0
$\theta^{2,*}$	0	0	0	0.0160	0	0	0	0	0	0.9761
$\theta^{3,*}$	0	0	-0.0165	0	0	0	0	-1.9777	0	0

TABLE 6

Errors of the solution x_1, x_2, x_3 of the system identified using PDHG compared to the true solution of the K-O system. The reference is obtained by numerically solving the true system (5.6) using a central finite difference scheme. These errors indicate that the errors in the system identification in Table 5 may be due to a lack of unique identifiability of the system using the given data points.

	x_1	x_2	x_3
Relative L^2 error (%)	0.2144	0.3718	0.3152

to the corresponding multitime HJ PDE at $(0, \lambda_1, \dots, \lambda_n)$. In our numerical experiments, we generate data points for training and testing by solving (5.6) numerically for $x_1, x_2, x_3 \in [0, 10]$ using MATLAB *ode45* [35] and then using a central finite difference scheme to approximate the time derivatives. We set $\lambda_j = 1 \forall j$ and $\gamma_k = 0.1 \forall k$.

Instead of the sparse regression techniques employed by SINDy, we use PDHG to minimize (5.7) (see section 4.4). Each iteration of PDHG involves minimizing a loss function of the form (5.7), but with ℓ_2 -regularization instead of ℓ_1 . Hence, this subproblem can be solved using our Riccati-based methods. As discussed in section 4.4, note that we only need to apply RK4 for the first iteration of PDHG, and every subsequent iteration can be solved using a change of bias. As such, we do not suffer from any error accumulation related to repeated applications of RK4. In Table 5, we see that we do indeed recover a sparse identification of the dynamics. However, we also incorrectly identify nonzero coefficients for the basis functions x_2 and x_3 . We note that this misidentification may be the result of a lack of unique identifiability of the system from the data points sampled. In fact, Table 6 shows that the errors in the solution x_1, x_2, x_3 of the identified system versus the solution of the true system (5.6) are relatively small, which corroborates that identifiability may have been an issue.

6. Summary. In this paper, we established a novel theoretical connection between regularized learning problems and the multitime Hopf formula. In doing so, we showed that when we solve these learning problems, we actually solve certain multitime HJ PDEs and their corresponding optimal control problems. In this work, we focused on the development of the connection between regularized linear regression and the LQR problem. By leveraging this connection, we developed a new methodology based on solving Riccati ODEs that allows us to design new training approaches for certain machine learning applications, such as continual learning, posttraining calibration, hyperparameter tuning and exploration of the associated Pareto front, and sparse dynamics identification. We also showed that our Riccati-based approach yields some promising computational advantages over conventional learning methods; after the original training, the models learned using our approach can be continually updated without having to retrain the entire model or having access to all of the

previous data, which could be particularly useful in continual learning [39, 28, 47] and federated learning [30, 26]. We note that our connection shares some similarities to the approach of NeuralODEs [10], which also connects ODEs to machine learning; however, NeuralODEs focuses on establishing new neural networks architectures using ODEs, whereas we make connections between (Riccati) ODEs (that arise from optimal control theory) and the training/optimization process of existing machine learning methods and neural networks. Our connection is also related to [45], which connects LQR problems with Kalman filters; however, [45] does not make any connections to PDEs or learning problems.

Thus, our theoretical connection and Riccati-based methodology present many exciting opportunities. Some possible future directions are as follows. While our Riccati-based methodology allows us to alter the hyperparameters and data points used in the learning problem without having to retrain on all previous data, it also requires that the original training be done using our Riccati-based methodology. It would allow for increased versatility if we could more easily combine our Riccati-based approach with other training methods. Additionally, in sections 4.4 and 5.4, we showed that our Riccati-based approach allows computations to be reused when using nonquadratic regularizations. However, in this case, the training process still had to be restarted if the hyperparameters, dataset, or regularization type is changed. It would allow for more flexibility if we could develop more adaptive processes for changing these aspects of the learning problem when the regularization is not quadratic.

In section 3, we focused on LQR problems, where the dynamics are independent of the trajectory, but it would be worthwhile to investigate what connections LQR problems with general linear dynamics may yield (e.g., LQR problems with general linear dynamics may be reformulated as LQR problems with state-independent dynamics using a change of variable [17]). Another natural extension would be to consider nonlinear models, which currently pose challenges in both scientific machine learning and optimal control and, hence, connections drawn in this case would benefit both fields. It would also be interesting to explore applications of our connection between general learning problems and general optimal control problems (i.e., the connection in section 2.2). In general, these optimal control problems can be tackled by solving characteristic line equations, which are represented by two-point boundary value problems (TPBVPs) and thus are more challenging to solve than initial value problems like the Riccati equation. However, working with these TPBVPs may be bypassed in some cases by considering numerical solvers such as those in [8, 9, 12, 15, 17, 50].

Alternatively, if we relax our assumption on convexity by allowing for nonconvex loss functions (or, equivalently, nonconvex Hamiltonians), we could also extend our connection to be between learning problems and differential games [19] instead of optimal control. However, even with convex losses, there are already many interesting potential applications in machine learning to pursue. For example, in the regularized linear regression case, it may be of interest to explore adjusting methods such as recursive least squares to perform hyperparameter estimation, e.g., similarly to [11]. More generally, the theoretical connection presented in section 2.2 for general convex losses provides a formula for the optimal learning parameters θ^* in terms of the hyperparameters. This connection could be leveraged to simplify the bilevel optimization in metalearning applications to a single-level, constrained optimization problem. Finally, we could generalize our theory to instead consider viscous HJ PDEs, which add a Laplacian term to the right-hand side of the HJ PDEs (2.1)

and (2.4). Then, by leveraging the recently established connection between viscous HJ PDEs and Bayesian modeling [14], we could extend to applications in Bayesian inference.

Appendix A. Glossary of related terms in each problem.

TABLE 7

Glossary of the mathematical relations between the quantities in the regularized learning problem (2.6), the multitime HJ PDE (2.4), and the associated optimal control problem (2.3). These mathematical relations allow us to form our theoretical connection in section 2.2. Note that the variables \mathbf{x}, t in the HJ PDE become hyperparameters in the learning problem, and we can treat them as constants when optimizing the learning problem (2.6) with respect to $\boldsymbol{\theta} = \mathbf{p}$.

Learning Problem	Multitime HJ PDE	Optimal control	Relation
Minimal loss function $\min_{\boldsymbol{\theta} \in \mathbb{R}^n} \mathcal{L}(\boldsymbol{\theta})$	Solution $S(\mathbf{x}, t_1, \dots, t_N)$	Value function $S(\mathbf{x}, t_1, \dots, t_N)$	$\min_{\boldsymbol{\theta} \in \mathbb{R}^n} \mathcal{L}(\boldsymbol{\theta}) = -S(\mathbf{x}, t_1, \dots, t_N)$
Optimal parameters $\boldsymbol{\theta}^*$	spatial gradient $\nabla_{\mathbf{x}} S(\mathbf{x}, t_1, \dots, t_N)$; maximizer \mathbf{p}^* in the Hopf formula	Optimal control $\mathbf{u}^*(\cdot)$	$\boldsymbol{\theta}^* = \mathbf{p}^* = \nabla_{\mathbf{x}} S(\mathbf{x}, t_1, \dots, t_N)$; $\mathbf{u}^*(s) \in \operatorname{argmax}_{\mathbf{u} \in \mathbb{R}^m} \{ \langle -f_i(\mathbf{u}), \mathbf{p}^* \rangle - L_i(\mathbf{u}) \}$, $s \in \left(\sum_{j=1}^{i-1} t_j, \sum_{j=1}^i t_j \right]$
Data fitting weights λ_i	Time t_i	Piecewise time horizons T_j	$\lambda_i = t_i, i = 1, \dots, N$; $T_j = \sum_{i=1}^j t_i, j = 0, \dots, N$
Data fitting loss $\mathcal{L}_i(\mathcal{AF}(\mathbf{z}_i; \cdot), \mathbf{y}_i)$	Hamiltonian $H_i(\cdot)$	Running cost $L_i(\mathbf{u})$, dynamics $f_i(\mathbf{u})$	$\mathcal{L}_i(\mathcal{AF}(\mathbf{z}_i; \cdot), \mathbf{y}_i) = H_i(\cdot) = \sup_{\mathbf{u} \in \mathbb{R}^m} \{ \langle -f_i(\mathbf{u}), \cdot \rangle - L_i(\mathbf{u}) \}$
Regularization $R(\cdot)$	Initial data $J(\cdot)$, spatial variable \mathbf{x}	Terminal cost $J(\cdot)$, initial position \mathbf{x}	$R(\mathbf{p}) = J^*(\mathbf{p}) - \langle \mathbf{x}, \mathbf{p} \rangle + c(\mathbf{x})$, where $c(\mathbf{x})$ is a constant (possibly 0) that may depend on \mathbf{x}
# of learnable parameters n	Spatial dimension n	Dimension of state space n	Same number n
# of data points N	# of times N	# of time intervals N	Same number N
Dimension of measured data m (e.g., $\mathcal{AF}(\cdot; \boldsymbol{\theta}), \mathbf{y}_i \in \mathbb{R}^m$)	Hamiltonian $H_i(\cdot)$	Dimension of control space m	$H_i(\cdot) = \sup_{\mathbf{u} \in \mathbb{R}^m} \{ \langle -f_i(\mathbf{u}), \cdot \rangle - L_i(\mathbf{u}) \}$

Appendix B. A more general version of LQR. In section 3.1, we presented a concise overview of the LQR problem in its canonical form. However, this paper establishes a linkage between regression problems and LQR problems with lower-order terms. Thus, in this section, we discuss an LQR problem with lower-order terms that more closely aligns with those used to establish our connections. We note that the notation employed in this section differs slightly from the notation utilized in the main body of the paper. The finite-horizon, continuous-time LQR with lower order terms can be expressed as follows,

$$(B.1) \quad S(\mathbf{x}, t) = \min_{\mathbf{u}(\cdot)} \left\{ \int_0^t \left(\frac{1}{2} \mathbf{x}(s)^T \mathcal{Q} \mathbf{x}(s) + \frac{1}{2} \mathbf{u}(s)^T \mathcal{R} \mathbf{u}(s) + \mathbf{x}(s)^T \mathcal{N} \mathbf{u}(s) - \mathbf{a}_u^T \mathbf{u}(s) - \mathbf{a}_x^T \mathbf{x}(s) \right) ds + \frac{1}{2} \mathbf{x}(t)^T \mathcal{Q}_f \mathbf{x}(t) + \mathbf{b}^T \mathbf{x}(t) : \dot{\mathbf{x}}(s) = \mathbf{A} \mathbf{x}(s) + \mathbf{B} \mathbf{u}(s) \forall s \in (0, t], \mathbf{x}(0) = \mathbf{x} \right\},$$

where $\mathcal{Q}, \mathcal{Q}_f \in \mathbb{R}^{n \times n}$ and $\mathcal{R} \in \mathbb{R}^{m \times m}$ are symmetric positive definite, $\mathcal{N} \in \mathbb{R}^{n \times m}$, $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, $\mathbf{a}_x, \mathbf{b} \in \mathbb{R}^n$, and $\mathbf{a}_u \in \mathbb{R}^m$. The corresponding HJ PDE is

$$(B.2) \quad \begin{cases} \frac{\partial S(\mathbf{x}, t)}{\partial t} + H(\mathbf{x}, \nabla_{\mathbf{x}} S(\mathbf{x}, t)) = 0, & \mathbf{x} \in \mathbb{R}^n, t > 0, \\ S(\mathbf{x}, 0) = J(\mathbf{x}), & \mathbf{x} \in \mathbb{R}^n, \end{cases}$$

where the initial data of the HJ PDE are given by the terminal cost $J(\mathbf{x}) := \frac{1}{2} \mathbf{x}^T \mathcal{Q}_f \mathbf{x} + \mathbf{b}^T \mathbf{x}$ of the optimal control problem (B.1) and the Hamiltonian H is defined by

$$(B.3) \quad \begin{aligned} H(\mathbf{x}, \mathbf{p}) &= \sup_{\mathbf{u} \in \mathbb{R}^m} \langle -f(\mathbf{x}, \mathbf{u}), \mathbf{p} \rangle - L(\mathbf{x}, \mathbf{u}) \\ &= -\langle A\mathbf{x}, \mathbf{p} \rangle - \frac{1}{2} \langle \mathbf{x}, \mathcal{Q}\mathbf{x} \rangle + \mathbf{a}_x^T \mathbf{x} + \frac{1}{2} \langle B^T \mathbf{p} + \mathcal{N}^T \mathbf{x} - \mathbf{a}_u, \mathcal{R}^{-1} (B^T \mathbf{p} + \mathcal{N}^T \mathbf{x} - \mathbf{a}_u) \rangle, \end{aligned}$$

where $f(\mathbf{x}, \mathbf{u}) = A\mathbf{x} + B\mathbf{u}$ is the source term of the dynamics and $L(\mathbf{x}, \mathbf{u}) = \frac{1}{2} \mathbf{x}^T \mathcal{Q}\mathbf{x} + \frac{1}{2} \mathbf{u}^T \mathcal{R}\mathbf{u} + \mathbf{x}^T \mathcal{N}\mathbf{u} - \mathbf{a}_u^T \mathbf{u} - \mathbf{a}_x^T \mathbf{x}$ is the running cost.

It is well known that this LQR problem can be solved using the Riccati equation as follows. Define $C_{pp} = B\mathcal{R}^{-1}B^T$, $C_{xx} = -\mathcal{N}\mathcal{R}^{-1}\mathcal{N}^T + \mathcal{Q}$, $C_{xp} = A - B\mathcal{R}^{-1}\mathcal{N}^T$, $\mathbf{v}_x = \mathbf{a}_x - \mathcal{N}\mathcal{R}^{-1}\mathbf{a}_u$, and $\mathbf{v}_p = B\mathcal{R}^{-1}\mathbf{a}_u$. Then, the solution is given by $S(\mathbf{x}, t) = \frac{1}{2} \mathbf{x}^T P(t)\mathbf{x} + \mathbf{q}(t)^T \mathbf{x} + r(t)$, where the function $P: [0, \infty) \rightarrow \mathbb{R}^{n \times n}$ which takes values in the space of symmetric positive definite matrices, $\mathbf{q}: [0, \infty) \rightarrow \mathbb{R}^n$, and $r: [0, \infty) \rightarrow \mathbb{R}$ solve the following Riccati ODE system:

$$(B.4) \quad \begin{cases} \dot{P}(t) = -P(t)^T C_{pp} P(t) + P(t)^T C_{xp} + C_{xp}^T P(t) + C_{xx}, & t \in (0, +\infty), \\ \dot{\mathbf{q}}(t) = -P(t)^T C_{pp} \mathbf{q}(t) + C_{xp}^T \mathbf{q}(t) - \mathbf{v}_x + P(t)^T \mathbf{v}_p, & t \in (0, +\infty), \\ \dot{r}(t) = -\frac{1}{2} \|\mathcal{R}^{-1/2} (B^T \mathbf{q}(t) - \mathbf{a}_u)\|^2, & t \in (0, +\infty), \\ P(0) = \mathcal{Q}_f, \mathbf{q}(0) = \mathbf{b}, r(0) = 0. \end{cases}$$

Appendix C. Details of the methodology.

C.1. Algorithm for deleting one data point. Here, we provide details for the algorithm for deleting one data point from section 4.2. Removing the j th data point corresponds to removing the term $\frac{1}{2} \lambda_j \|\Phi_j \boldsymbol{\theta} - \mathbf{y}_j\|_2^2$ in the loss function (4.1) or, equivalently, removing the Hamiltonian $\frac{1}{2} \|\Phi_j \boldsymbol{\theta} - \mathbf{y}_j\|_2^2$ from the multitime HJ PDE and removing the pieces $L_j(s, \mathbf{u}) = \frac{1}{2} \mathbf{u}^T \mathbf{u} - \mathbf{y}_j$ and $f(s, \mathbf{u}) = \Phi_j^T \mathbf{u}$ from the running cost and dynamics, respectively, of the corresponding piecewise LQR problem. Hence, numerically, we can remove the j th data point by solving the following Riccati ODE,

$$(C.1) \quad \begin{cases} \dot{\tilde{P}}(t) = -\tilde{P}(t)^T \Phi_j^T \Phi_j \tilde{P}(t), & t < \lambda_j, \\ \dot{\tilde{\mathbf{q}}}(t) = -\tilde{P}(t)^T \Phi_j^T (\Phi_j \tilde{\mathbf{q}}(t) - \mathbf{y}_j), & t < \lambda_j \end{cases}$$

with terminal condition $\tilde{P}(\lambda_j) = P(T_N)$ and $\tilde{\mathbf{q}}(\lambda_j) = \mathbf{q}(T_N)$, where $P(T_N)$ and $\mathbf{q}(T_N)$ are obtained from solving the learning problem (4.1) with all N data points. Then, the solution to the new learning problem with the j th point removed is given by (4.4), where $\tilde{P} = \tilde{P}(0)$ and $\tilde{\mathbf{q}} = \tilde{\mathbf{q}}(0)$ are the solution to (C.1).

C.2. Algorithm for tuning the regularization weights. Here, we provide details for the algorithm for tuning the regularization weights from section 4.3. We consider the case where we change each regularization parameter γ_k to $\tilde{\gamma}_k$. This change can be regarded as two steps: first, we change all parameters γ_k for the indices k such that $\tilde{\gamma}_k > \gamma_k$, and then we change the other parameters. Define the index set \mathcal{K} to be $\mathcal{K} = \{k: \tilde{\gamma}_k > \gamma_k\}$.

The first step is equivalent to adding the term $\sum_{k \in \mathcal{K}} \frac{\tilde{\gamma}_k - \gamma_k}{2} (\theta_k - \theta_k^0)^2$ to the loss function (4.1). We can interpret this as adding an $(N + 1)$ th Hamiltonian $\frac{1}{2} \boldsymbol{\theta}^T \Gamma_+ \boldsymbol{\theta}$ with corresponding time variable $t_{N+1} = 1$ to the multitime HJ PDE, where Γ_+ is a diagonal matrix whose k th diagonal element is $\tilde{\gamma}_k - \gamma_k$ if $k \in \mathcal{K}$ and 0 otherwise. Therefore, the solution to this new multitime HJ PDE can be solved by the following Riccati equation,

$$(C.2) \quad \begin{cases} \dot{P}_+(t) = -P_+(t)^T \Gamma_+ P_+(t), & t \in (0, 1), \\ \dot{\mathbf{q}}_+(t) = -P_+(t)^T \Gamma_+ \mathbf{q}_+(t), & t \in (0, 1) \end{cases}$$

with initial condition $P_+(0)$ and $\mathbf{q}_+(0)$, which are the corresponding solutions to the Riccati equations before changing the weights $\gamma_k, k \in \mathcal{K}$. In other words, we set $P_+(0) = P(T_N)$ and $\mathbf{q}_+(0) = \mathbf{q}(T_N)$, where $P(T_N), \mathbf{q}(T_N)$ are obtained from solving the original learning problem (4.1) with the original values of γ_k .

The second step is equivalent to removing the term $\sum_{k \notin \mathcal{K}} \frac{\gamma_k - \tilde{\gamma}_k}{2} (\theta_k - \theta_k^0)^2$ from the loss function (4.1). This is equivalent to solving a single-time HJ PDE with a terminal condition at time 1 and Hamiltonian $\frac{1}{2} \boldsymbol{\theta}^T \Gamma_- \boldsymbol{\theta}$, where Γ_- is a diagonal matrix whose k th diagonal element is $\gamma_k - \tilde{\gamma}_k$ if $k \notin \mathcal{K}$ and 0 otherwise. Then, the solution can be obtained by solving the following Riccati equation,

$$(C.3) \quad \begin{cases} \dot{P}_-(t) = -P_-(t)^T \Gamma_- P_-(t), & t \in (0, 1), \\ \dot{\mathbf{q}}_-(t) = -P_-(t)^T \Gamma_- \mathbf{q}_-(t), & t \in (0, 1) \end{cases}$$

with terminal condition $P_-(1) = P_+(1)$ and $\mathbf{q}_-(1) = \mathbf{q}_+(1)$, where P_+, \mathbf{q}_+ are obtained from the solution to (C.2).

Finally, the minimizer of the new learning problem after changing all of the weights γ_k in the regularization term is given by $P_-(0) \tilde{\Gamma} \boldsymbol{\theta}^0 + \mathbf{q}_-(0)$, where P_-, \mathbf{q}_- are obtained from the solution to (C.3) and $\tilde{\Gamma}$ is a diagonal matrix whose k th diagonal element is the new regularization parameter $\tilde{\gamma}_k$.

Appendix D. Computational complexity of the methodology. In this section, we compare the computational complexity of our Riccati-based methodology (using RK4) from section 4 to that of the method of least squares.

D.1. Initial training. First, we consider solving the learning problem (4.1) with N points, which may be considered as an initial training step with N data points.

In this case, our Riccati-based methodology requires us to solve a sequence of N Riccati ODEs of the form (4.2). Using RK4 to solve each ODE involves computing $\Phi_i^T \Phi_i$ (recall that $\Phi_i \in \mathbb{R}^{m \times n}$), which requires $O(mn^2)$ operations, and multiplying several $n \times n$ matrices together, which requires $O(n^3)$ operations. In other words, solving one step of Riccati ODEs (i.e., adding one data point to our training set) requires $O(mn^2 + n^3)$ operations. Thus, the overall computational complexity of our Riccati-based methodology for training on N points is $O(Nmn^2 + Nn^3)$.

In contrast, using the method of least squares to solve (4.1) with N points involves computing $\Phi_i^T \Phi_i$ for each point, which requires $O(mn^2)$ operations for each point or $O(Nmn^2)$ operations for all N points, and then solving an $n \times n$ linear system, which requires $O(n^3)$ operations. Thus, the overall computational complexity of the method of least squares for training on N points is $O(Nmn^2 + n^3)$.

As expected, the method of least squares beats our Riccati-based methodology in this case, but as we discussed previously, our methodology is advantageous when changing the training set or hyperparameters.

D.2. Continual learning. In this section, we discuss adding data points to the training set (as in section 4.2), where we update our learned models as soon as a new data point becomes available (e.g., as in continual learning). We consider the initial training set to have N points and then sequentially incorporate K more data points.

As discussed in section D.1, adding one data point using our Riccati-based methodology requires $O(mn^2 + n^3)$ operations. Thus, the overall computational complexity of using our methodology to perform the initial training and then to retrain K times is $O((N+K)mn^2 + (N+K)n^3)$. Adding the i th data point using the method of least squares requires $O((N+i)mn^2 + n^3)$ operations. Thus, the overall computational complexity of using the method of least squares to perform the initial training and then to retrain K times is $\sum_{i=0}^K O((N+i)mn^2 + n^3) = O((KN + K^2)mn^2 + Kn^3)$. Thus, our Riccati-based methodology becomes increasingly more advantageous over the method of least squares as the number of points added, K , increases.

D.3. Tuning the bias. As discussed in section 4.3, tuning the bias θ^0 using our Riccati-based methodology is accomplished via the update (4.5). This update only involves a matrix-vector multiplication (recall that Γ is diagonal) and vector addition. Thus, tuning the bias using our methodology requires $O(n^2)$ operations. In contrast, tuning the bias using the method of least squares requires retraining on the entire dataset. Thus, if the training set has size N , then the computational complexity of the method of least squares in this case is $O(Nmn^2 + n^3)$. Hence, our methodology is much more efficient than the method of least squares in this case. Note that these same calculations also apply to using our methodology versus the method of least squares to perform the first step in each iteration of PDHG (4.7), which, as discussed previously, is equivalent to performing a change of bias.

Appendix E. Riccati ODEs and recursive least squares. In section 4.2, we discuss adding or removing data using our Riccati-based methodology. However, data points could also be added or removed incrementally using recursive least squares. In this section, we discuss the connections between our Riccati-based methodology and recursive least squares. Namely, we show that the recursive least squares method for this problem can be considered as a special case of our Riccati-based methodology. Note that in this section, we focus on adding data points, but the results for removing data points can be derived similarly. For generality, we consider the more general learning problem (3.15) in section 3.3, and the results for the more specific learning problem in section 4.2 can be recovered using the identifications listed in section 4.1.

The analytical solution to the Riccati ODEs (3.12), (3.13) is given by

$$P_N = \left(\mathcal{Q}_f^{-1} + \sum_{i=1}^N t_i B_i \mathcal{R}_i^{-1} B_i^T \right)^{-1},$$

$$\mathbf{q}_N = \left(\mathcal{Q}_f^{-1} + \sum_{i=1}^N t_i B_i \mathcal{R}_i^{-1} B_i^T \right)^{-1} \left(\mathcal{Q}_f^{-1} \mathbf{b} + \sum_{i=1}^N t_i B_i \mathcal{R}_i^{-1} \mathbf{a}_i \right),$$

where $P_N = P \left(\sum_{i=1}^N t_i \right)$ and $\mathbf{q}_N = \mathbf{q} \left(\sum_{i=1}^N t_i \right)$. Then, using the Woodbury matrix identity, the recursive least squares solution is given by

(E.1)

$$P_N = P_{N-1} - t_N P_{N-1} B_N (\mathcal{R}_N + t_N B_N^T P_{N-1} B_N)^{-1} B_N^T P_{N-1},$$

$$\mathbf{q}_N = \mathbf{q}_{N-1} + t_N P_N B_N \mathcal{R}_N^{-1} \mathbf{a}_N - t_N P_{N-1} B_N (\mathcal{R}_N + t_N B_N^T P_{N-1} B_N)^{-1} B_N^T \mathbf{q}_{N-1}.$$

PROPOSITION E.1. *The recursive least squares solution (E.1) is a discretization of the Riccati ODEs (3.12), (3.13).*

Proof. We begin by proving the result for P . Let $P(\cdot)$ be the solution to the Riccati ODEs. Then $P(s)$ is invertible for all $s \in [0, \sum_{i=1}^N t_i]$, which gives us that $\frac{dP^{-1}}{ds} = -P^{-1} \frac{dP}{ds} P^{-1}$. Plugging this identity into (3.12), we have that

$$(E.2) \quad \frac{dP^{-1}(s)}{ds} = B_i \mathcal{R}_i^{-1} B_i^T, \quad s \in [T_{i-1}, T_i], \quad P^{-1}(0) = Q_f^{-1}.$$

The solution to this ODE is given by

$$P^{-1}(s) = Q_f^{-1} + \sum_{j=1}^i t_j B_j \mathcal{R}_j^{-1} B_j^T + (s - T_{i-1}) B_i \mathcal{R}_i^{-1} B_i^T \quad \forall s \in [T_{i-1}, T_i], i = 0, \dots, N.$$

If we only consider the values of $P^{-1}(s)$ at the discrete times $s = T_0, T_1, \dots, T_N$, then this solution can be defined recursively as

$$P^{-1}(T_i) = P^{-1}(T_{i-1}) + t_i B_i \mathcal{R}_i^{-1} B_i^T$$

or, in other words, we have that

$$P(T_i) = (P^{-1}(T_{i-1}) + t_i B_i \mathcal{R}_i^{-1} B_i^T)^{-1}.$$

By applying the Woodbury matrix identity, we recover the recursive least squares solution for P in (E.1). Next, we prove the result for \mathbf{q} . Consider

$$\frac{d(P^{-1}\mathbf{q})}{ds}(s) = \frac{dP^{-1}}{ds} \mathbf{q} + P^{-1} \frac{d\mathbf{q}}{ds} = B_i \mathcal{R}_i^{-1} B_i^T \mathbf{q} + P^{-1} \frac{d\mathbf{q}}{ds} = B_i \mathcal{R}_i^{-1} \mathbf{a}_i$$

for $s \in [T_{i-1}, T_i]$ and with initial condition $(P^{-1}\mathbf{q})(0) = -Q_f^{-1}\mathbf{b}$, where the second equality follows from (E.2) and the third equality follows from multiplying (3.13) by P^{-1} on the left. Then, the solution to this ODE gives us that

$$\mathbf{q}(s) = P(s)P^{-1}(T_{i-1})\mathbf{q}(T_{i-1}) + (s - T_{i-1})P(s)B_i \mathcal{R}_i^{-1} \mathbf{a}_i \quad \forall s \in [T_{i-1}, T_i], i = 0, \dots, N,$$

which defines $\mathbf{q}(s)$, recursively. If we only consider the values of $\mathbf{q}(s)$ at the discrete times $s = T_0, T_1, \dots, T_N$, then using this recursive definition of \mathbf{q} simplifies to

$$\mathbf{q}(T_i) = P(T_i) (P^{-1}(T_{i-1})\mathbf{q}(T_{i-1}) + t_i B_i \mathcal{R}_i^{-1} \mathbf{a}_i).$$

Finally, applying the recursive least squares solution for $P(T_i)$ in (E.1) to the definition of $\mathbf{q}(T_i)$ above gives us the recursive least squares solution for \mathbf{q} in (E.1). \square

Thus, we have shown that the recursive least squares solution represents a special case (in the form of a discretization) of the Riccati ODEs and, hence, we may regard the recursive least squares update (E.1) as an alternative method for solving the Riccati ODEs (3.12), (3.13), if we only care about the solution at the discrete times $s = T_0, \dots, T_N$. For instance, in section 5.3, we show how using more general numerical methods like RK4 naturally expose 1D curves along the Pareto front of solutions, which aligns with the Riccati ODEs as a continuous formulation; in contrast, the discretization used by the recursive least squares update would only allow for discrete points on the Pareto front. Moreover, we note that there are some cases in which recursive least squares is known to be numerically unstable [34, 31, 32], whereas Runge–Kutta methods have well-established stability and accuracy results [6].

Appendix F. Additional results for example 2. In section 5.2, we consider a 1D steady-state reaction-diffusion equation (5.2) and discuss three different types of posttraining calibrations: adding new data points to compensate for a lack of knowledge in the regular training, enforcing the fitting of some data points by increasing the weights λ_i of their respective terms in the loss function (5.3), and removing some data points so that their effects are eliminated. In this section, we present the results for the last case and then provide a large-scale continual learning example for solving (5.2). In Figure 9, we remove two outliers one-by-one and observe that their removal does successfully increase the accuracy of the learned model. Again, using our Riccati-based approach, the removal of these points is done using only knowledge about the point to be removed and the results of the previous training step.

Next, we provide a large-scale continual learning example for solving the 1D steady-state reaction-diffusion equation (5.2). Similarly to section 5.1, the measurements of the source term f are corrupted by additive Gaussian noise with a relatively large noise scale. Hence, we need a large amount of data to solve (5.2) accurately. We again follow the continual learning setup and assume that the measurements of f must be accessed in a stream, such that the historical data are not available after they are incorporated into our learned model. Specifically, each new data point of f comes from uniformly randomly sampling a point $x \in [0, 1]$ and is corrupted by additive Gaussian noise with noise scale two. The same basis functions and loss function listed in section 5.2 are used in this case. The results are presented in Table 8 and Figure 10. We observe that the accuracy of our inferences of u and f is improved continuously as new measurements of f are incorporated. This process does not re-

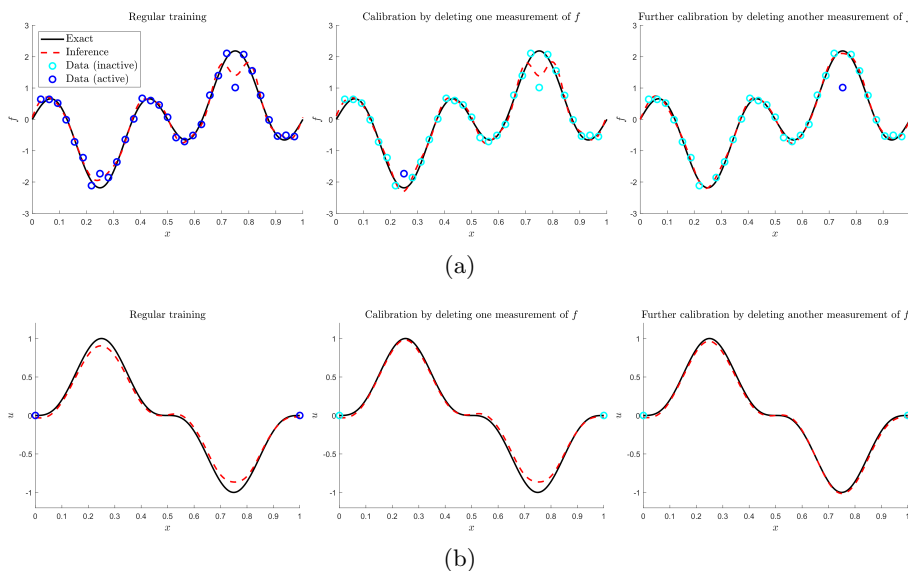


FIG. 9. Results of solving the 1D steady-state reaction-diffusion equation (5.2) with noisy measurements of the source term f in the domain and noiseless measurements of the solution u on the boundary. (a) results for f ; (b) results for u . Left: results of regular training using our Riccati-based method in section 4.1; middle and right: calibrating the results of regular training by eliminating two outlier measurements of f using the methodology described in section 4.2. Calibrations are employed without retraining or access to the data from the previous training. Note: color appears only in the online article.

TABLE 8

Errors of u and f when solving the 1D steady-state reaction-diffusion equation (5.2) with different numbers N of noisy measurements of f using our Riccati-based approach. The errors represent the relative L^2 errors on a uniform grid of $[0, 1]$. New measurements are incorporated following the continual learning framework, and, using our Riccati-based methodology, incorporating new data points does not require retraining on or access to any of the previous data points. Qualitative results can be found in Figure 10.

	$N = 100$	$N = 1000$	$N = 50000$
Error of u	72.59%	22.55%	1.76%
Error of f	89.65%	29.76%	3.35%

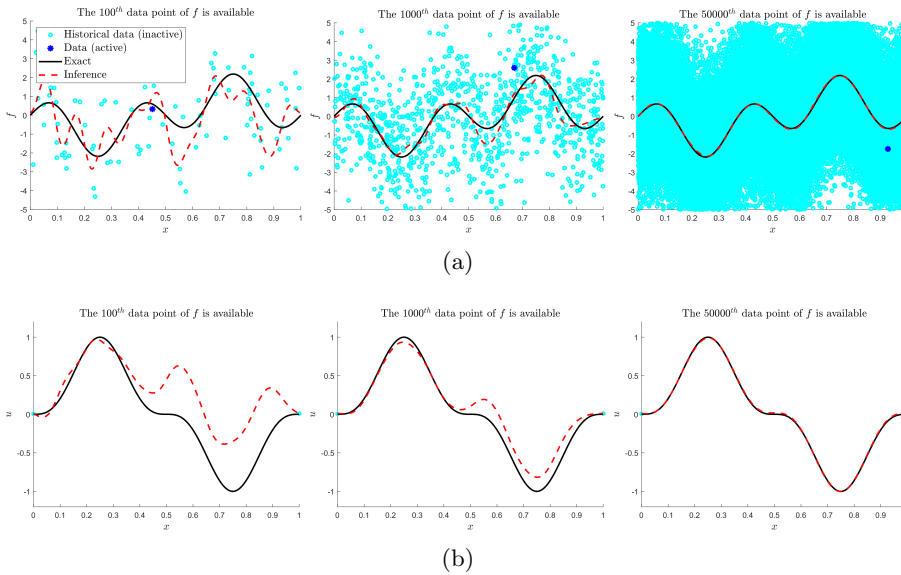


FIG. 10. Continual learning of the source term f and solution u to the 1D steady-state reaction-diffusion equation (5.2) of our Riccati-based approach as more data become available. Measurements of the source term f are corrupted by additive Gaussian noise with noise scale two. Our Riccati-based approach allows us to incrementally update the learned coefficients as more data become available without requiring access to the previous data or retraining on the entire dataset, which provides advantages in both memory and computations over conventional learning methods. Note: color appears only in the online article.

quire storage or use of historical data, which again highlights the computational and memory advantages of our Riccati-based approach.

Appendix G. Details of the hyperparameters in the numerical examples. In this section, we include additional details of the numerical examples from section 5. Recall that RK4 is used to numerically solve the associated Riccati ODEs. The step size of RK4, denoted as h , is chosen based on three principles: high accuracy, efficient computations, and avoiding computational overflow. Unless otherwise stated, we use Python, the NumPy library [23], the TensorFlow library [1], and double precision in all of the numerical examples.

In the first example (section 5.1), $\lambda_i = 1 \forall i$, $\gamma_k = 0.1 \forall k$, and the step size h of RK4 is set to be 0.001, 0.0005, 0.0001. A uniform grid with size 1001 on $[0, 10]$ is used for evaluation (computing relative L^2 error). In the second example (section 5.2), we change λ_b from 1 to 10 and the rest of the hyperparameters remain the same as the

TABLE 9

Step size h of RK4 when tuning γ in the hyperparameter tuning/Pareto front visualization example in section 5.3. $\gamma_0 \rightarrow \gamma_1$ represents the case when tuning γ from γ_0 to γ_1 .

γ	$1 \rightarrow 10^{-1}$	$10^{-1} \rightarrow 10^{-2}$	$10^{-2} \rightarrow 10^{-3}$	$10^{-3} \rightarrow 10^{-4}$	$10^{-4} \rightarrow 10^{-5}$
h	10^{-2}	10^{-3}	10^{-4}	10^{-6}	10^{-7}

regular training. The time step of RK4 is $h = 0.0001$. A uniform grid with size 257 on $[0, 1]$ is used for evaluation. In the third example (section 5.3), the multihead PINN used to obtain the basis functions has three hidden layers, each of which is equipped with hyperbolic tangent activation functions and 100 neurons. The multihead PINN is trained with the Adam optimizer [27], where the learning rate is 1×10^{-3} and the other hyperparameters are set to their default TensorFlow values. The training data for f are evaluated on a 33×33 uniform grid on $[0, 1]^2$ and the inference of u is evaluated on a 257×257 uniform grid. The results of the regular training are obtained using $\lambda_i = 1 \forall i$, $\gamma_k = 1 \forall k$, and $h = 0.001$. When tuning γ and visualizing the Pareto front, we use the values of γ and the step size h of RK4 listed in Table 9. In the last example (section 5.4), we set $\lambda_i = 1 \forall i$, $\gamma_k = 0.1 \forall k$, and $h = 0.001$. In our implementation of the PDHG method, we set $\sigma_\theta = \sigma_w = 0.5$. The solution x_1, x_2, x_3 of the system identified is evaluated on a uniform grid of 10,001 on $[0, 10]$.

REFERENCES

- [1] M. ABADI, A. AGARWAL, P. BARHAM, *TensorFlow: Large-scale Machine Learning on Heterogeneous Systems*, <https://www.tensorflow.org/> (2015).
- [2] B. D. ANDERSON AND J. B. MOORE, *Optimal Control: Linear Quadratic Methods*, Courier Corporation, Mineola, NY, 2007.
- [3] M. BARDI AND I. CAPUZZO-DOLCETTA, *Optimal Control and Viscosity Solutions of Hamilton-Jacobi-Bellman Equations*, Mod. Birkhäuser Class., Birkhäuser Boston, Boston, MA, 1997, <https://doi.org/10.1007/978-0-8176-4755-1>.
- [4] M. BARDI AND L. EVANS, *On Hopf's formulas for solutions of Hamilton-Jacobi equations*, *Nonlinear Anal.*, 8 (1984), pp. 1373–381, [https://doi.org/10.1016/0362-546X\(84\)90020-8](https://doi.org/10.1016/0362-546X(84)90020-8).
- [5] S. L. BRUNTON, J. L. PROCTOR, AND J. N. KUTZ, *Discovering governing equations from data by sparse identification of nonlinear dynamical systems*, *Proc. Natl. Acad. Sci. USA*, 113 (2016), pp. 3932–3937.
- [6] J. C. BUTCHER, *Numerical Methods for Ordinary Differential Equations*, Wiley, Chichester, England, 2016.
- [7] A. CHAMBOLLE AND T. POCK, *A first-order primal-dual algorithm for convex problems with applications to imaging*, *J. Math. Imaging Vision*, 40 (2011), pp. 120–145.
- [8] P. CHEN, J. DARBON, AND T. MENG, *Hopf-type representation formulas and efficient algorithms for certain high-dimensional optimal control problems*, *Comput. Math. Appl.*, to appear.
- [9] P. CHEN, J. DARBON, AND T. MENG, *Lax-Oleinik-Type formulas and efficient algorithms for certain high-dimensional optimal control problems*, *Commun. Appl. Math. Comput.*, to appear.
- [10] R. T. CHEN, Y. RUBANOVA, J. BETTENCOURT, AND D. K. DUVENAU, *Neural ordinary differential equations*, *Adv. Neural Inf. Process. Syst.* 31, Curran Associates, Red Hook, NY, 2018, pp. 6572–6583.
- [11] J. CHUNG, M. CHUNG, J. T. SLAGEL, AND L. TENORIO, *Sampled limited memory methods for massive linear inverse problems*, *Inverse Problems*, 36 (2020), 054001.
- [12] J. DARBON, *On convex finite-dimensional variational methods in imaging sciences and Hamilton-Jacobi equations*, *SIAM J. Imaging Sci.*, 8 (2015), pp. 2268–2293, <https://doi.org/10.1137/130944163>.
- [13] J. DARBON, P. M. DOWER, AND T. MENG, *Neural network architectures using min-plus algebra for solving certain high-dimensional optimal control problems and Hamilton-Jacobi PDEs*, *Math. Control Signals Systems*, 35 (2023), pp. 1–44, <https://doi.org/10.1007/s00498-022-00333-2>.
- [14] J. DARBON AND G. LANGLOIS, *On Bayesian posterior mean estimators in imaging sciences and Hamilton-Jacobi partial differential equations*, *J. Math Imaging Vision*, 63 (2021), pp. 821–854.

- [15] J. DARBON AND T. MENG, *On decomposition models in imaging sciences and multi-time Hamilton–Jacobi partial differential equations*, SIAM J. Imaging Sci., 13 (2020), pp. 971–1014, <https://doi.org/10.1137/19M1266332>.
- [16] J. DARBON, T. MENG, AND E. RESMERITA, *On Hamilton–Jacobi PDEs and image denoising models with certain nonadditive noise*, J. Math. Imaging Vision, 64 (2022), pp. 408–441.
- [17] J. DARBON AND S. OSHER, *Algorithms for overcoming the curse of dimensionality for certain Hamilton–Jacobi equations arising in control theory and elsewhere*, Res. Math. Sci., 3 (2016), pp. 1–26, <https://doi.org/10.1186/s40687-016-0068-7>.
- [18] S. DESAI, M. MATTHEAKIS, H. JOY, P. PROTOPAPAS, AND S. ROBERTS, *One-shot transfer learning of physics-informed neural networks*, in the 2nd AI4Science Workshop at the 39th International Conference on Machine Learning (ICML), Baltimore, MD, 2022.
- [19] L. C. EVANS AND P. E. SOUGANIDIS, *Differential games and representation formulas for solutions of Hamilton–Jacobi–Isaacs equations*, Indiana Univ. Math. J., 33 (1984), pp. 773–797.
- [20] I. GOODFELLOW, Y. BENGIO, AND A. COURVILLE, *Deep Learning*, MIT Press, Cambridge, MA, 2016.
- [21] S. GOSWAMI, K. KONTOLATI, M. D. SHIELDS, AND G. E. KARNIADAKIS, *Deep transfer operator learning for partial differential equations under conditional shift*, Nature Mach. Intell., 4 (2022), pp. 1155–1164.
- [22] J. HAN, A. JENTZEN, AND W. E, *Solving high-dimensional partial differential equations using deep learning*, Proc. Natl. Acad. Sci. USA, 115 (2018), pp. 8505–8510.
- [23] C. R. HARRIS ET AL., *Array programming with NumPy*, Nature, 585 (2020), pp. 357–362.
- [24] E. HOPF, *Generalized solutions of non-linear equations of first order*, J. Math. Mech., 14 (1965), pp. 951–973.
- [25] Y. JIN AND B. SENDHOFF, *Pareto-based multiobjective machine learning: An overview and case studies*, IEEE Trans. Syst. Man Cybernet., Part C, 38 (2008), pp. 397–415, <https://doi.org/10.1109/TSMCC.2008.919172>.
- [26] P. KAIROUZ ET AL., *Advances and open problems in federated learning*, Found. Trends Mach. Learn., 14 (2021), pp. 1–210.
- [27] D. P. KINGMA AND J. BA, *Adam: A method for stochastic optimization*, in Proceedings of the 3rd International Conference on Learning Representations (ICLR 2015), San Diego, CA, 2015, Conference Track Proceedings, 2015.
- [28] J. KIRKPATRICK ET AL., *Overcoming catastrophic forgetting in neural networks*, Proc. Natl. Acad. Sci. USA, 114 (2017), pp. 3521–3526.
- [29] Y. LECUN, Y. BENGIO, AND G. HINTON, *Deep learning*, Nature, 521 (2015), pp. 436–444.
- [30] T. LI, A. K. SAHU, A. TALWALKAR, AND V. SMITH, *Federated learning: Challenges, methods, and future directions*, IEEE Signal Process. Mag., 37 (2020), pp. 50–60.
- [31] A. P. LIAVAS AND P. A. REGALIA, *Numerical stability issues of the conventional recursive least squares algorithm*, in Proceedings of the 1998 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP’98, Vol. 3, IEEE, Piscataway, NJ, 1998, pp. 1409–1412.
- [32] A. P. LIAVAS AND P. A. REGALIA, *On the numerical stability and accuracy of the conventional recursive least squares algorithm*, IEEE Trans. Signal Process., 47 (1999), pp. 88–96.
- [33] P. L. LIONS AND J.-C. ROCHET, *Hopf formula and multitime Hamilton–Jacobi equations*, Proc. Amer. Math. Soc., 96 (1986), pp. 79–84, <http://www.jstor.org/stable/2045657>.
- [34] S. LJUNG AND L. LJUNG, *Error propagation properties of recursive least-squares adaptation algorithms*, Automatica, 21 (1985), pp. 157–167.
- [35] MathWorks Inc., *MATLAB for Artificial Intelligence*, <https://www.mathworks.com>.
- [36] W. MCENEANEY, *Max-plus Methods for Nonlinear Control and Estimation*, Birkhäuser, Boston, 2006.
- [37] M. MOHRI, A. ROSTAMIZADEH, AND A. TALWALKAR, *Foundations of Machine Learning*, MIT Press, Cambridge, MA, 2018.
- [38] T. NAKAMURA-ZIMMERER, Q. GONG, AND W. KANG, *QRnet: Optimal regulator design with LQR-augmented neural networks*, IEEE Control Syst. Lett., 5 (2021), pp. 1303–1308.
- [39] G. I. PARISI, R. KEMKER, J. L. PART, C. KANAN, AND S. WERMTER, *Continual lifelong learning with neural networks: A review*, Neural Netw., 113 (2019), pp. 54–71.
- [40] A. F. PSAROS, X. MENG, Z. ZOU, L. GUO, AND G. E. KARNIADAKIS, *Uncertainty quantification in scientific machine learning: Methods, metrics, and comparisons*, J. Comput. Phys., 477 (2023), 111902.

- [41] M. RAISSI, P. PERDIKARIS, AND G. KARNIADAKIS, *Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations*, J. Comput. Phys., 378 (2019), pp. 686–707.
- [42] J.-C. ROCHET, *The taxation principle and multi-time Hamilton-Jacobi equations*, J. Math. Econom., 14 (1985), pp. 113–128.
- [43] S. J. RUSSELL, *Artificial Intelligence: A Modern Approach*, Pearson Education, 2010.
- [44] J. SIRIGNANO AND K. SPILIOPOULOS, *DGM: A deep learning algorithm for solving partial differential equations*, J. Comput. Phys., 375 (2018), pp. 1339–1364.
- [45] E. TODOROV, *General duality between optimal control and estimation*, in 2008 47th IEEE Conference on Decision and Control, IEEE, Piscataway, NJ, 2008, pp. 4286–4292.
- [46] H. L. TRENTELMAN, A. A. STOOORVOGEL, AND M. HAUTUS, *Control Theory for Linear Systems*, Springer, London, 2001.
- [47] G. M. VAN DE VEN AND A. S. TOLIAS, *Three Scenarios for Continual Learning*, *NeurIPS Continual Learning Workshop*, 1.9 (2018).
- [48] X. WAN AND G. E. KARNIADAKIS, *Multi-element generalized polynomial chaos for arbitrary probability measures*, SIAM J. Sci. Comput., 28 (2006), pp. 901–928.
- [49] S. WEISBERG, *Applied Linear Regression*, Wiley Ser. Probab. Stat. 528, Wiley, Hoboken, NJ, 2005.
- [50] I. YEGOROV AND P. M. DOWER, *Perspectives on characteristics based curse-of-dimensionality-free numerical approaches for solving Hamilton–Jacobi equations*, Appl. Math. Optim., 83 (2021), pp. 1–49.
- [51] Z. ZHANG, Z. ZOU, E. KUHL, AND G. E. KARNIADAKIS, *Discovering a reaction-diffusion model for Alzheimer’s disease by combining PINNs with symbolic regression*, Comput. Methods Appl. Mech. Engrg., 419 (2024), 116647.
- [52] Z. ZOU AND G. E. KARNIADAKIS, *L-HYDRA: Multi-Head Physics-Informed Neural Networks*, preprint, arXiv:2301.02152, 2023.
- [53] Z. ZOU, X. MENG, A. F. PSAROS, AND G. E. KARNIADAKIS, *NeuralUQ: A comprehensive library for uncertainty quantification in neural differential equations and operators*, SIAM Rev., 66 (2024), pp. 161–190, <https://doi.org/10.1137/22M1518189>.