

PDE Control Gym: A Benchmark for Data-Driven Boundary Control of Partial Differential Equations

Luke Bhan*

University of California, San Diego

LBHAN@UCSD.EDU

Yuexin Bian*

University of California, San Diego

YUBIAN@UCSD.EDU

Miroslav Krstic

University of California, San Diego

KRSTIC@UCSD.EDU

Yuanyuan Shi

University of California, San Diego

YYSHI@UCSD.EDU

Editors: A. Abate, M. Cannon, K. Margellos, A. Papachristodoulou

Abstract

Over the last decade, data-driven methods have surged in popularity, emerging as valuable tools for control theory. As such, neural network approximations of control feedback laws, system dynamics, and even Lyapunov functions have attracted growing attention. With the ascent of learning based control, the need for accurate, fast, and easy-to-use benchmarks has increased. In this work, we present the first learning-based environment for boundary control of PDEs. In our benchmark, we introduce three foundational PDE problems - a 1D transport PDE, a 1D reaction-diffusion PDE, and a 2D Navier–Stokes PDE - whose solvers are bundled in an user-friendly reinforcement learning gym. With this gym, we then present the first set of model-free, reinforcement learning algorithms for solving this series of benchmark problems, achieving stability, although at a higher cost compared to model-based PDE backstepping. With the set of benchmark environments and detailed examples, this work significantly lowers the barrier to entry for learning-based PDE control - a topic largely unexplored by the data-driven control community. The entire benchmark is available on [Github](#) along with detailed [documentation](#) and the presented reinforcement learning models are [open sourced](#).

Keywords: Partial Differential Equation Control, Nonlinear Systems, Benchmarking for Data-Driven Control, Reinforcement Learning

1. Introduction

As learning-based control has exploded across both academia and industry, the need for fast, and accurate bench-marking is heightened. For example, perhaps the most visible impact of proper bench-marking is in the field of computer-vision resulting in 15-years of breakthrough results from AlexNet [Krizhevsky et al. \(2012\)](#) to neural radiance fields (NeRFs) [Mildenhall et al. \(2020\)](#). Despite this, the control community has, justifiably, forgone consistent efforts in bench-marking as the community spawned from an applied mathematics perspective where the focus was behind *provable* stability guarantees. However, given the recent exploration surrounding data-driven control methods [Berberich et al. \(2023\)](#); [Feng et al. \(2023\)](#), designing fast, well-documented, and challenging benchmarks is of utmost importance to ensure new learning-based control approaches are consistently advancing the state of the art.

* equal contribution

In this work, we develop a benchmarking suite for learning-based boundary control of PDEs. Boundary control of PDEs is of elevated importance compared to control across the full domain as many real world problems *cannot* control the PDE across the entire domain, but only at the boundary input. Thus, boundary control is physically more realistic as the actuation and sensing are generally non-intrusive [Krstic and Smyshlyaev \(2008b\)](#). For example, in fluid flows, the engineer only gets control access to the surrounding walls containing the fluid [Di Meglio et al. \(2012\)](#) or in temperature manufacturing, the engineer is typically unable to set the temperature of entire plate, but only a specific edge. Furthermore, boundary control is extremely powerful in modeling macro-level traffic congestion [Huan Yu \(2023\)](#) as modern highways typically only enable actuation and sensing at the on/off ramps. However, for a majority of these applications, each researcher typically develops their own simulations and thus there is no standard library with a universal set of problems to test new algorithms. Thus, in this work, we introduce the first library containing a set of general PDE control problems and implementations of their corresponding model-based control algorithms that can be easily modified to fit the wide array of aforementioned target applications.

Contributions This paper has three main contributions. First, we introduce, design, and formalize the first benchmark suite for PDE control including 3 classical problems ranging from boundary stabilization for 1D transport (hyperbolic) and reaction-diffusion (parabolic) PDEs to trajectory following for the 2D Navier-Stokes PDEs. Along with the proposal of the PDE control benchmarking suite, we parameterize the numerical scheme implementations as RL gyms - effectively decoupling the PDE solvers from the controller design, enabling the use of *any pre-implemented learning algorithm* for PDE control. Second, utilizing our benchmark suite, we train the first set of *model-free* RL controllers which effectively stabilize hyperbolic and parabolic PDE problems and achieve effective tracking for the 2D Navier-Stokes equations. We then compare the resulting controllers to classical algorithms such as PDE backstepping and adjoint-based optimization highlighting the trade-offs between performance. Lastly, we provide extensive documentation and numerous examples for the training of RL controllers, implementation of classical control algorithms and of course the integration of new PDE control problems into the benchmark suite.

2. Related Work

2.1. Learning-based PDE benchmarks

Benchmarking for machine learning in PDEs	Compilation of a premade dataset	Supports control	Differentiable PDE solver	Supports custom PDEs	Supports reinforcement learning	Implementation of model-based control
PhiFlow Holl et al. (2020)	✗	✓	✓	✓	✗	✗
PDEBench Takamoto et al. (2022) (created from PhiFlow)	✓	✗	✗	✗	✗	✗
PDEArena Gupta and Brandstetter (2022) (created from PhiFlow)	✓	✗	✗	✗	✗	✗
PDE Control Gym (Ours)	✗	✓	✗	✓	✓	✓

Table 1: Comparison of benchmarks for machine learning in PDEs.

Currently, to the author’s knowledge, there are no benchmarking suites focused on *boundary control* of PDEs as most benchmarks such as [Takamoto et al. \(2022\)](#); [Gupta and Brandstetter \(2022\)](#) present datasets for learning PDE solution maps from initial conditions. These benchmarks are typically used for comparing neural network-based PDE solvers like neural operators [Lu et al. \(2021\)](#);

Li et al. (2021) and PINNs Raissi et al. (2019). Although these benchmarks are effective, they do not allow users to incorporate boundary control or alter the PDEs within their datasets. The ϕ_{Flow} framework Holl et al. (2020) provides a flexible PDE solver that supports automatic differentiation for the calculation PDE derivatives to be used in control algorithms. While adaptable for PDE control, it does not natively support RL algorithms and lacks a set of model-based controllers for learning-based control comparisons. Lastly, it is worth noting that while there are ODE control suites like Duan et al. (2016); Tassa et al. (2018) tailored for ODE control tasks. The PDE control gym, to our knowledge, represents the *first* PDE-focused benchmarking suite for learning-based control algorithms.

2.2. Learning as a tool for PDE control

As with most scientific disciplines, machine learning has had a broad impact in PDE control. In 1D PDE problems, a series of work has been developed to use neural operators for approximating control feedback laws Bhan et al. (2023a,b); Krstic et al. (2024); Qi et al. (2023), under a supervised learning framework with provable stability guarantees. Furthermore, an optimal control approach using PINNs is explored in Mowlavi and Nabi (2023). Additionally, the closest paper to this work is by Yu et al. (2022) who presented the first exploration utilizing RL for PDE boundary control. However, they do not explore the benchmark PDE problems presented in this paper instead focusing on Aw-Rascale-Zhang (ARZ) traffic model.

2.3. Reinforcement learning in control

Reinforcement learning (RL) has demonstrated significant success in various control applications, including robotics Brunke et al. (2022), power systems Chen et al. (2022), and autonomous driving Kiran et al. (2021). From a controls perspective, deep RL (DRL) algorithms learn feedback laws that maps observations (states) into actions (control inputs), typically via neural networks (NN). These RL controllers are trained to optimize specific reward functions, such as the L^2 spatial norm of states for stabilization tasks Yu et al. (2022). The most appealing feature of deep RL is its *model-free* nature, allowing it to control complex systems without requiring explicit model estimations. Consequently, RL has potential to outperform model-based control methods in highly complex tasks with hard-to-model dynamics. To demonstrate the use of PDE Control Gym, we conduct experiments using off-the-shelf RL algorithms implemented with Stable-Baselines3 Raffin et al. (2021). We selected the off-policy soft actor-critic (SAC) (Haarnoja et al., 2018) and on-policy proximal policy optimization (PPO) (Schulman et al., 2017) algorithms for their demonstrated efficiency in solving challenging continuous control tasks Duan et al. (2016).

3. Formalization of PDE Control Problems

3.1. General PDE control problem

We consider a partial differential equation (PDE) defined on a domain \mathcal{X} , which can be either one-dimensional (1D), $\mathcal{X} = [0, 1] \subset \mathbb{R}$, or two-dimensional (2D), $\mathcal{X} = [0, 1] \times [0, 1] \subset \mathbb{R}^2$. The time domain is $\mathcal{T} = [0, T] \subset \mathbb{R}^+$. Let $u(x, t)$, $x \in \mathcal{X}$, $t \in \mathcal{T}$ describe the state of the system governed by the PDE according to the dynamics

$$\frac{\partial u}{\partial t} = \mathcal{P} \left(u, \frac{\partial u}{\partial x}, \frac{\partial^2 u}{\partial x^2}, \dots, U(t) \right), \quad (1)$$

Supported Configuration	1D Hyperbolic		1D Parabolic		2D Navier-Stokes	
	Sensing ($o(t)$)	Actuation ($a(t)$)	Sensing ($o(t)$)	Actuation ($a(t)$)	Sensing ($o(t)$)	Actuation ($a(t)$)
full-state	$u(x, t)$	$u(1, t)$	$u(x, t)$	$u(1, t)$	$u(x, y, t)$	$u(x, 1, t)$
collocated	$u_x(1, t)$	$u(1, t)$	$u_x(1, t)$	$u(1, t)$	—	—
anti	$u(0, t)$	$u(1, t)$	—	—	—	—
collocated	$u_x(0, t)$	$u(1, t)$	$u_x(0, t)$	$u(1, t)$	—	—
full-state	$u(x, t)$	$u_x(1, t)$	$u(x, t)$	$u_x(1, t)$	$u(x, y, t)$	$u(x, 0, t)$
collocated	$u(1, t)$	$u_x(1, t)$	$u(1, t)$	$u_x(1, t)$	—	—
anti	$u(0, t)$	$u_x(1, t)$	—	—	—	—
collocated	$u_x(0, t)$	$u_x(1, t)$	$u_x(0, t)$	$u_x(1, t)$	—	—
full-state	—	—	—	—	$u(x, y, t)$	$u(1, y, t)$
full-state	—	—	—	—	$u(x, y, t)$	$u(0, y, t)$

Table 2: Configurations for actuation and sensing supported by the PDE Control Gym for the three problems. Full state indicates the measurement is the entire PDE state, collocated indicates that the sensing and measurement is done at the same boundary point, and anti-collocated indicates sensing and measurement are done at opposite boundary points. The configurations marked in blue correspond to the experiment examples in Section 5.

where \mathcal{P} is the partial differential equation(s) that model(s) the system dynamics, and $U(t)$ is the control function. Then, the goal of a PDE control problem is to optimize a cost function (e.g. regulate $u(x, t)$ to be the desired trajectory while reducing the control cost, stabilize the PDEs) from just boundary inputs. Note that in some methods such as PDE backstepping, optimization is forgone in favor of just asymptotic stabilization as the infinite dimensional nature of PDE control is extremely challenging.

3.2. Markov decision processes (MDPs) for PDE control

We give a brief overview of the components of the MDP governing both the 1D and 2D support problems. More details about the specific MDPs governing the examples in Section 5 can be found in the [supplemental](#).

State and Observation Space. In all PDE control problems addressed, the state $s(t)$ at time t is represented by the PDE value $u(x, t)$, $x \in \mathcal{X}$. To enhance flexibility for different PDE tasks such as observer design, we have developed different partial state measurement settings, which we denote as the observation space $o(t)$. The types of sensing supported for each problem are detailed in Table 3.2. Furthermore, we offer the ability to introduce custom noise functions (with built-in support for Gaussian noise). This allows users to simulate real-world sensor noise in their experiments.

Action Space. The action $a(t) = U(t)$ for both the RL and control agents is determined by the actuation locations and boundary condition type. In 1D hyperbolic and parabolic systems, we consider both Neumann and Dirichlet boundary actuation $U(t) \in \mathbb{R}$ at either boundary, with four possible cases: 1) $u(0, t) = U(t)$, 2) $u(1, t) = U(t)$, 3) $u_x(0, t) = U(t)$, and 4) $u_x(1, t) = U(t)$.

Considering the symmetry between the boundaries $x = 0$ and $x = 1$, there are eight distinct combinations for the 1D Hyperbolic PDE problem and six for the 1D Parabolic PDE problem, including an additional boundary condition at $u(0, t)$. These combinations are outlined in the first two sections of Table 3.2. For the 2D Navier–Stokes problem, we consider Dirichlet-type boundary actuation on any of the four boundaries: top, bottom, left, and right. The gym also allows users to customize actuator positions, enabling research into optimizing both location and actuation type.

State Evolution. To simulate the PDE system evolution with the control input $U(t)$, we use a first-order Taylor approximation for temporal evolution,

$$u(t + \Delta t) = u(t) + \Delta t \cdot \mathcal{P} \left(u(t), \frac{\partial u}{\partial x}, \frac{\partial^2 u}{\partial x^2}, \dots, U(t) \right). \quad (2)$$

Spatial derivatives are approximated using appropriate finite difference schemes and are explicitly given for each gym environment in the [supplemental](#). In practice, selecting the time step Δt and spatial discretization Δx for each problem requires careful consideration, particularly based on the number of approximated spatial derivatives. Nonetheless, we found that reasonable choices such as $\Delta x = 0.01$, $\Delta t = 0.0001$ yield both fast and numerically stable results.

Reward. Reward shaping plays a pivotal role in the training of RL algorithms. Generally speaking, for stabilization tasks, it is appropriate to employ a form of trajectory-based reward

$$\int_0^T \left(\int_{x \in \mathcal{X}} \|u(x, t)\|^2 dx + \|U(t)\|^2 \right) dt, \quad (3)$$

which minimizes the state magnitudes and control efforts. For tracking tasks, a trajectory reward as

$$\int_0^T \left(\int_{x \in \mathcal{X}} \|u(x, t) - u_{ref}(x, t)\|^2 dx + \|U(t) - U_{ref}(t)\|^2 \right) dt, \quad (4)$$

is a reasonable choice as it penalizes deviations from the reference trajectory given in both state $u_{ref}(x, t)$ and control actions $U_{ref}(t)$. However, in practice, for the 1D hyperbolic and parabolic PDE problems, we found that the reward as given in (3) was insufficient for training and thus we use a specifically tuned reward that penalizes the difference of the L_2 norms between the current state and next state after action $a(t)$ (presented in [supplemental](#)).

4. Benchmark PDE Control Tasks

4.1. 1D Hyperbolic (transport) PDEs

We consider the benchmark transport PDE in the form

$$u_t(x, t) = u_x(x, t) + \beta(x)u(0, t), \quad (5)$$

for $x \in [0, 1], t \in [0, T]$. Physically, (5) is a “transport process (from $x = 1$ towards $x = 0$) with recirculation” of the outlet variable $u(0, t)$. Recirculation causes *instability* and the goal is to stabilize “full-state” recirculation from only boundary inputs. In practice, we consider the same PDE as studied in [Bhan et al. \(2023a\)](#) where β is governed by the Chebyshev polynomial $\beta(x) = 5 \cos(\gamma \cos^{-1}(x))$ and Dirichlet actuation $u(1, t) = U(t)$. Classically, this PDE, with recirculation, has been a seminal-benchmark for PDE backstepping, as the 1D transport PDE can model a wide range of applications from chemical processes to shallow water waves and traffic flows [Krstic and Smyshlyaev \(2008a\)](#).

Model-based backstepping control. The backstepping controller given by the following

$$U(t) = \int_0^1 k(1-y)u(y,t)dy, \quad (6)$$

$$k(x) = -\beta(x) + \int_0^x \beta(x-y)k(y)dy, \quad (7)$$

for $x \in [0, 1]$. (6) results in stabilization of (5) Krstic and Smyshlyaev (2008a). In practice, the backstepping kernel (7) is implemented using the successive approximations approach although a Laplace transform approach is also viable Krstic and Smyshlyaev (2008b).

4.2. 1D Parabolic (reaction-diffusion) PDEs

We consider the benchmark reaction-diffusion PDEs governed by recirculation function $\lambda(x)$ as

$$u_t(x,t) = u_{xx}(x,t) + \lambda(x)u(x,t), \quad (8)$$

$$u(0,t) = 0, \quad (9)$$

with Dirichlet or Neumann actuation at $x = 1$. Again, instability is caused by the $\lambda(x)u(x,t)$ term otherwise the problem would simplify to the classical heat equation. This PDE appears in different applications ranging from a chemical tubular reactor Shi et al. (2022) to electro-chemical battery models Moura et al. (2014) and diffusion in social networks Wang et al. (2020).

Model-based backstepping control. For the PDE (8), (9), with Dirichlet boundary actuation $u(1,t) = U(t)$, the backstepping controller with full state measurement is given by the following Smyshlyaev and Krstic (2004, 2010),

$$U(t) = \int_0^1 k(1,y)u(y,t)dy. \quad (10)$$

where $k(x,y) \in C^2(\tilde{\mathcal{T}})$, $\tilde{\mathcal{T}} = \{0 \leq y \leq x \leq 1\}$.

$$k_{xx}(x,y) - k_{yy}(x,y) = \lambda(y)k(x,y), \quad \forall (x,y) \in \check{\mathcal{T}}, \quad (11)$$

$$k(x,0) = 0, \quad (12)$$

$$k(x,x) = -\frac{1}{2} \int_0^x \lambda(y)dy, \quad (13)$$

where $\check{\mathcal{T}} = \{0 < y \leq x < 1\}$.

4.3. 2D Navier-Stokes PDEs

We consider the 2D in-compressible Navier-Stokes equations as the third benchmark control task,

$$\nabla \cdot \mathbf{u} = 0, \quad (14a)$$

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\frac{1}{\rho} \nabla p + \nu \nabla^2 \mathbf{u} \quad (14b)$$

With slight abuse of notation, we denote the spatial variable (in 2D) as $\mathbf{x} = (x,y) \in \mathcal{X} = [0, 1] \times [0, 1]$, and $\mathbf{u} = (u,v) : \mathcal{X} \times \mathcal{T} \rightarrow \mathbb{R}^2$ represents 2D velocity field, ν is the kinematic viscosity of

the fluid, ρ is the fluid density, and p is the pressure field. Navier-Stokes equation is fundamental in fluid dynamics with extensive applications including aerodynamic design, pollution modeling, and wind turbine flows [Jameson et al. \(1998\)](#); [Li et al. \(2021\)](#). For the experiments, we consider boundary control along the top boundary as $\mathbf{u}(x, 0, t) = U(x, t), \forall x \in [0, 1]$. All other boundary conditions are Dirichlet boundary conditions where the velocity is set to be 0. The task here is to find boundary control $U(x, t)$ such that the resulting velocity field is close to the reference trajectory given in both desired velocity field $\mathbf{u}_{ref}(x, t)$ and desired actions $U_{ref}(x, t)$.

Model-based optimization-based control. We provide an optimization-based controller based on [Pyta et al. \(2015\)](#) as the model-based control baseline.

$$\begin{aligned} \min_{U(x,t)} J(U(\cdot, t), \mathbf{u}) &= \frac{1}{2} \int_{\mathcal{T}} \int_{\mathcal{X}} \|\mathbf{u}(x, t) - \mathbf{u}_{ref}(x, t)\|^2 dx dt \\ &+ \frac{\gamma}{2} \int_{\mathcal{T}} \|U(\cdot, t) - U_{ref}(\cdot, t)\|^2 dt \end{aligned} \tag{15a}$$

$$\text{s.t.} \quad (14a), (14b), \quad \mathbf{u}(x, 0, t) = U(x, t), \forall x \in [0, 1]. \tag{15b}$$

To ensure computational tractability, the control actions are set to be the tangential, uniform velocity, i.e., $u(x, 0, t) = U(t), v(x, 0, t) = 0$, followed in [Pyta et al. \(2015\)](#). The optimal control actions are obtained by solving the PDE-constrained optimization problem presented in (15). This solution employs Lagrange multipliers, utilizing the adjoint method [McNamara et al. \(2004\)](#); [Gunzburger \(2002\)](#) for gradient computation of the Lagrangian function.

5. Experiments

For each of the 3 environments in PDE Control Gym, we implemented baseline model-based control algorithms as well as off-the-shelf RL algorithms including soft actor-critic (SAC) ([Haarnoja et al., 2018](#)) and proximal policy optimization (PPO) ([Schulman et al., 2017](#)) trained using Stable-Baselines3 (Parameters available in [supplemental](#)). We note that all the experiments can be trained in under 1 hour (Nvidia RTX 3090ti) and entire trajectories can be simulated in seconds.

Algorithm	Hyperbolic PDE	Parabolic PDE	Navier-Stokes PDE
	Average Episode Reward for Trained Policy \uparrow	Average Episode Reward for Trained Policy \uparrow	Average Episode Reward for Trained Policy \uparrow
Model-based	246.3	299.1	-7.931
PPO	172.3	293.3	-5.370
SAC	184.2	229.1	-17.829

Table 3: Resulting control algorithm performance on 50 test episodes in each gym (larger value indicates better performance). The model-based algorithm for the hyperbolic and parabolic PDEs are the backstepping schemes given in (6), (7) and (10), (11), (12), (13) respectively while the method for the Navier-Stokes PDE solves the optimization problem (15).




Figure 1: Rewards for training PPO (blue) and SAC (orange) on the 1D transport PDE, 1D reaction-diffusion PDE, and 2D Navier-Stokes PDE from left to right. The solid lines represent the mean and the shaded bounds are 95% confidence intervals across 5 seeds.

5.1. 1D Hyperbolic (transport) PDEs




Figure 2: Example of the 1D transport PDE system stabilization using backstepping, PPO, and SAC (left to right) under initial conditions $u(x, 0) = 10$. The recirculation coefficient is defined as $\beta(x) = 5 \cos(\gamma \cos^{-1}(x))$ with $\gamma = 7.35$.

Experimental design Our experimental setup for the Hyperbolic 1D problem, detailed in Section 4.1, considers full state measurements and boundary control $u(1, t) = U(t)$. We use the Chebyshev polynomial recirculation function $\beta(x) = 5 \cos(\gamma \cos^{-1}(x))$ from Bhan et al. (2023a), with $\gamma = 7.35$ (future studies may vary γ). Each episode is initiated from a random initial condition, $u(x, 0) \sim \text{Uniform}(1, 10)$. This setup presents a challenging control scenario, as the open loop system ($U(t) = 0$) is unstable (See Figure 1 in [supplemental](#)).

Results We now present detailed results on the policies trained and their comparison with model-based backstepping. In the left of Figure 1, we present the average reward functions for both RL algorithms over 1 million training steps. Then, in Table 3, we present the average reward where we run the trained final RL policies, and the model-based backstepping policy for 50 test episodes with different initial conditions, noting that model-based backstepping performs the best. Additionally, in Figures 2 we provide a comparison across all 3 control approaches where $u(x, 0) = 10$. We can clearly see that although all 3 policies are stabilizing for the examples, model-based PDE backstepping again performs the best and the RL control signals are high oscillatory leaving room for improvement in applying model-free PDE control algorithms.

5.2. 1D Parabolic (reaction-diffusion) PDEs




Figure 3: Example of reaction-diffusion PDE system stabilization using backstepping, PPO, and SAC (left to right) under initial conditions $u(x, 0) = 10$. The recirculation coefficient using the Chebyshev polynomial defined as $\lambda(x) = 50 \cos(\gamma \cos^{-1}(x))$ with $\gamma = 8$.

Experimental design We adopt the same approach as the 1D Hyperbolic PDE in Section 5.1 except that the dynamics are now governed by (8), (9), with full state measurements and $u(1, t) = U(t)$. The time-horizon is shortened to 1 second as the algorithms are able to stabilize faster. We choose $\lambda(x) = 50 \cos(\gamma \cos^{-1}(x))$ where γ is fixed to be 8 (future studies may vary γ). At each episode, the initial conditions are uniformly randomized according to $u(x, 0) \sim \text{Uniform}(1, 10)$, and we note that the system is always open-loop unstable for all possible initial conditions (See Figure 5 supplemental). For training, we follow the same procedure as Section 5.1 except that we require a finer simulation resolution of $\Delta x = 0.005$, and the PDE is simulated at $\Delta t = 0.00001$ due to the approximation of the second spatial derivative in the reaction-diffusion PDE.

Results Figure 1 presents the average reward over 1 million training steps. Unlike the hyperbolic PDE where both RL algorithms performed relatively equal, the PPO algorithm achieved better performance during training which is corroborated by the testing rewards in the middle column of Table 3. Figure 3 demonstrates a test case with $u(x, 0) = 10$. Similar to the transport PDE, we observe oscillations in the RL feedback laws, suggesting potential improvement via enforcing continuity constraints as in Asadi et al. (2018). Notably, in Figure 3, perhaps due to reward shaping, PPO differs from the backstepping controller’s approach, but maintains excellent performance.

5.3. 2D Navier-Stokes PDEs

Experimental design For the Navier-Stokes 2D problem (Section 4.3), both velocity components are zero initially, i.e., $u(x, y, 0) = v(x, y, 0) = 0$. We apply boundary control on the top boundary with tangential, uniform controlled velocity, setting $u(x, 1, t) = U(t) \in \mathbb{R}$ and $v(x, 0, t) = 0$. For implementation, we discretize the state space with a spatial step of $\Delta x = 0.05$ and the PDE is simulated at $\Delta t = 0.001$. The reward for training is derived from the negative of the cost in optimization (15). The reference velocity vector \mathbf{u}_{ref} is the resulted velocity vector under the boundary control $U(t) = 3 - 5t$, and $U_{ref} = 2.0$.

Results Figure 1 (right) shows the average reward per episode for PPO and SAC, with PPO outperforming SAC both in terms of higher final rewards and more stable training curves. Table 3 (right) details the average episodic rewards over 50 test episodes, where PPO surpasses both SAC and the model-based optimization algorithm, which often gets stuck in local optima. Despite the reward difference, on a singular example presented in Figure 4, all methods effectively track the reference velocity vectors.

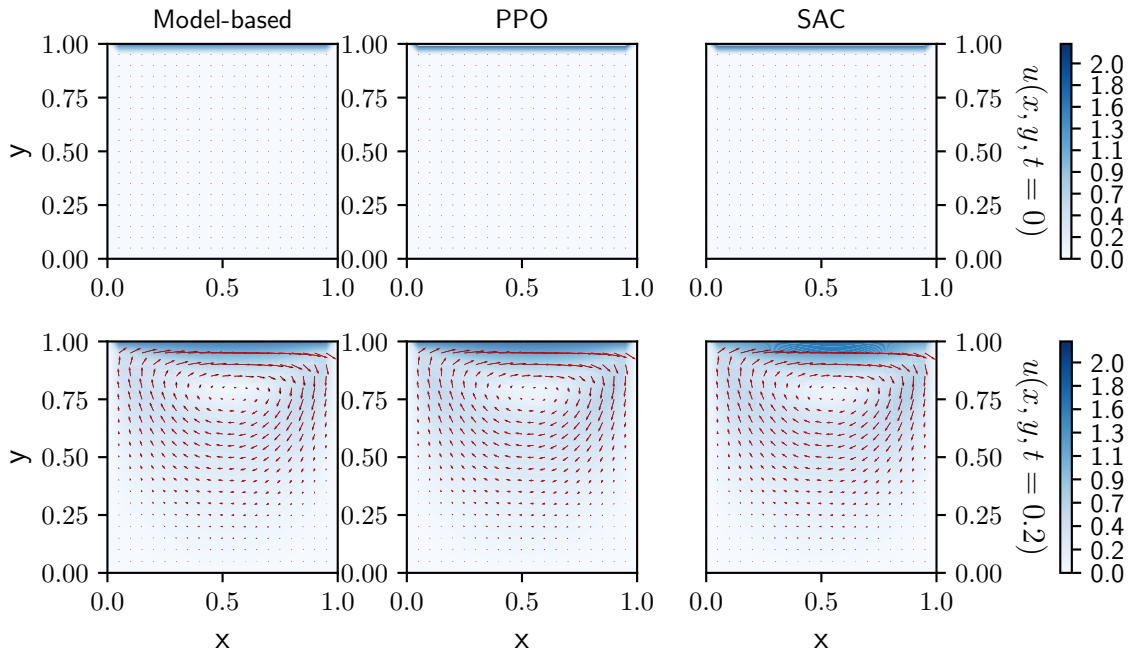


Figure 4: Example of Navier-Stokes PDE tracking using optimization-based control, PPO, and SAC under initial conditions $u(x, y, 0) = 0$ at $t = 0$ (top) and $t = 0.2$ (bottom). Red and black arrows represent the actual and reference velocity field respectively. The background color represents the magnitude of the velocity vector.

6. Conclusion

Future work Throughout this paper, we have mentioned several avenues for future research based on the PDE Control Gym. As such, we conclude this work by briefly summarizing these ideas. We employed relatively simple policy network architectures in our RL algorithms, not fully fine-tuning them to the specific problems. The PDE Control Gym presents opportunities to optimize policy network structures, improve reward shaping, and develop better RL algorithms for PDE control tasks. Additionally, our experiments were based on time-invariant linear instability coefficients where $\beta(x)$ and $\lambda(x)$ are unknown but static during RL training. Thus, there is much to be explored for model-free controllers when considering time-varying models, adaptive control, and sensing noise. Furthermore, given the superior performance of backstepping controllers, investigating the potential of pre-training RL methods through imitation learning could be a valuable direction.

Conclusion In this study, we introduced the first benchmark suite for learning-based boundary control of PDEs. We developed RL gyms for three fundamental PDE control problems: the 1D transport PDE, 1D reaction-diffusion PDE, and 2D Navier Stokes PDE. This gym allows for the separation of algorithm design from the numerical implementation of PDEs. Moreover, we trained a series of *model-free* RL models on the three benchmarks and compared their performance with model-based PDE backstepping and optimization methods. Finally, our work discussed multiple avenues for future research, aiming to inspire new research in the challenging field of PDE control.

References

- Kavosh Asadi, Dipendra Misra, and Michael L. Littman. Lipschitz continuity in model-based reinforcement learning, 2018.
- Julian Berberich, Carsten W. Scherer, and Frank Allgöwer. Combining prior knowledge and data for robust controller design. *IEEE Transactions on Automatic Control*, 68(8):4618–4633, 2023. doi: 10.1109/TAC.2022.3209342.
- Luke Bhan, Yuanyuan Shi, and Miroslav Krstic. Neural operators for bypassing gain and control computations in PDE backstepping. *IEEE Transactions on Automatic Control*, pages 1–16, 2023a. doi: 10.1109/TAC.2023.3347499.
- Luke Bhan, Yuanyuan Shi, and Miroslav Krstic. Operator learning for nonlinear adaptive control. In *Proceedings of The 5th Annual Learning for Dynamics and Control Conference (L4DC)*, volume 211 of *Proceedings of Machine Learning Research*, pages 346–357. PMLR, 15–16 Jun 2023b.
- Lukas Brunke, Melissa Greeff, Adam W Hall, Zhaocong Yuan, Siqi Zhou, Jacopo Panerati, and Angela P Schoellig. Safe learning in robotics: From learning-based control to safe reinforcement learning. *Annual Review of Control, Robotics, and Autonomous Systems*, 5:411–444, 2022.
- Xin Chen, Guannan Qu, Yujie Tang, Steven Low, and Na Li. Reinforcement learning for selective key applications in power systems: Recent advances and future challenges. *IEEE Transactions on Smart Grid*, 13(4):2935–2958, 2022.
- Florent Di Meglio, Rafael Vazquez, Miroslav Krstic, and Nicolas Petit. Backstepping stabilization of an underactuated 3×3 linear hyperbolic system of fluid flow equations. In *2012 American Control Conference (ACC)*, pages 3365–3370, 2012.
- Yan Duan, Xi Chen, Rein Houthoofd, John Schulman, and Pieter Abbeel. Benchmarking deep reinforcement learning for continuous control. In *International conference on machine learning*, pages 1329–1338. PMLR, 2016.
- Jie Feng, Yuanyuan Shi, Guannan Qu, Steven H. Low, Anima Anandkumar, and Adam Wierman. Stability constrained reinforcement learning for decentralized real-time voltage control. *IEEE Transactions on Control of Network Systems*, pages 1–12, 2023. doi: 10.1109/TCNS.2023.3338240.
- Max D Gunzburger. *Perspectives in flow control and optimization*. SIAM, 2002.
- Jayesh K Gupta and Johannes Brandstetter. Towards multi-spatiotemporal-scale generalized PDE modeling. arXiv:2209.15616, 2022.
- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pages 1861–1870. PMLR, 2018.
- Philipp Holl, Nils Thuerey, and Vladlen Koltun. Learning to control PDEs with differentiable physics. In *International Conference on Learning Representations (ICLR)*, 2020.

- Miroslav Krstic Huan Yu. *Traffic Congestion Control by PDE Backstepping*. Birkhäuser Cham, 2023.
- A. Jameson, L. Martinelli, and N. A. Pierce. Optimum aerodynamic design using the Navier-Stokes equations. *Theoretical and Computational Fluid Dynamics*, 10(1):213–237, Jan 1998.
- B Ravi Kiran, Ibrahim Sobh, Victor Talpaert, Patrick Mannion, Ahmad A Al Sallab, Senthil Yogamani, and Patrick Pérez. Deep reinforcement learning for autonomous driving: A survey. *IEEE Transactions on Intelligent Transportation Systems*, 23(6):4909–4926, 2021.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C.J. Burges, L. Bottou, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 25, 2012.
- Miroslav Krstic and Andrey Smyshlyaev. Backstepping boundary control for first-order hyperbolic PDEs and application to systems with actuator and sensor delays. *Systems & Control Letters*, 57(9):750–758, 2008a.
- Miroslav Krstic and Andrey Smyshlyaev. *Boundary Control of PDEs*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2008b.
- Miroslav Krstic, Luke Bhan, and Yuanyuan Shi. Neural operators of backstepping controller and observer gain functions for reaction–diffusion PDEs. *Automatica*, 164:111649, 2024. ISSN 0005-1098. doi: <https://doi.org/10.1016/j.automatica.2024.111649>. URL <https://www.sciencedirect.com/science/article/pii/S0005109824001420>.
- Zongyi Li, Nikola Borislavov Kovachki, Kamyar Azizzadenesheli, Burigede liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. In *International Conference on Learning Representations (ICLR)*, 2021.
- Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3(3):218–229, 2021.
- Antoine McNamara, Adrien Treuille, Zoran Popović, and Jos Stam. Fluid control using the adjoint method. *ACM Transactions On Graphics (TOG)*, 23(3):449–456, 2004.
- Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020.
- Scott J Moura, Nalin A Chaturvedi, and Miroslav Krstić. Adaptive partial differential equation observer for battery state-of-charge/state-of-health estimation via an electrochemical model. *Journal of Dynamic Systems, Measurement, and Control*, 136(1):011015, 2014.
- Saviz Mowlavi and Saleh Nabi. Optimal control of PDEs using physics-informed neural networks. *Journal of Computational Physics*, 473:111731, 2023.
- Lorenz Pyta, Michael Herty, and Dirk Abel. Optimal feedback control of the incompressible Navier-Stokes-equations using reduced order models. In *2015 54th IEEE Conference on Decision and Control (CDC)*, pages 2519–2524. IEEE, 2015.

- Jie Qi, Jing Zhang, and Miroslav Krstic. Neural operators for delay-compensating control of hyperbolic PIDEs. arXiv:2307.11436, 2023.
- Antonin Raffin, Ashley Hill, Adam Gleave, Anssi Kanervisto, Maximilian Ernestus, and Noah Dornmann. Stable-Baselines3: Reliable Reinforcement Learning Implementations. *Journal of Machine Learning Research*, 22(268):1–8, 2021.
- M. Raissi, P. Perdikaris, and G.E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. arXiv:1707.06347, 2017.
- Yuanyuan Shi, Zongyi Li, Huan Yu, Drew Steeves, Anima Anandkumar, and Miroslav Krstic. Machine learning accelerated PDE backstepping observers. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 5423–5428, 2022.
- A. Smyshlyaev and M. Krstic. Closed-form boundary state feedbacks for a class of 1-D partial integro-differential equations. *IEEE Transactions on Automatic Control*, 49(12):2185–2202, 2004.
- A. Smyshlyaev and M. Krstic. *Adaptive Control of Parabolic PDEs*. Princeton University Press, 2010.
- Makoto Takamoto, Timothy Praditia, Raphael Leiteritz, Daniel MacKinlay, Francesco Alesiani, Dirk Pflüger, and Mathias Niepert. PDEBench: an extensive benchmark for scientific machine learning. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 1596–1611, 2022.
- Yuval Tassa, Yotam Doron, Alistair Muldal, Tom Erez, Yazhe Li, Diego de Las Casas, David Budden, Abbas Abdolmaleki, Josh Merel, Andrew Lefrancq, Timothy P. Lillicrap, and Martin A. Riedmiller. Deepmind control suite. *CoRR*, abs/1801.00690, 2018.
- Haiyan Wang, Feng Wang, and Kuai Xu. *Modeling information diffusion in online social networks with partial differential equations*, volume 7. Springer Nature, 2020.
- Huan Yu, Saehong Park, Alexandre Bayen, Scott Moura, and Miroslav Krstic. Reinforcement learning versus PDE backstepping and PI control for congested freeway traffic. *IEEE Transactions on Control Systems Technology*, 30(4):1595–1611, 2022.