

# Spatio-Temporal Stochastic Optimization: Theory and Applications to Optimal Control and Co-Design

Ethan N. Evans, Andrew P. Kendall, George I. Boutselis, and Evangelos A. Theodorou  
Department of Aerospace Engineering, Georgia Institute of Technology  
Email: eevans41@gatech.edu

**Abstract**—There is a rising interest in Spatio-temporal systems described by Partial Differential Equations (PDEs) among the control community. Not only are these systems challenging to control, but the sizing and placement of their actuation is an NP-hard problem on its own. Recent methods either discretize the space before optimization, or apply tools from linear systems theory under restrictive linearity assumptions. In this work we consider control and actuator placement as a coupled optimization problem, and derive an optimization algorithm on Hilbert spaces for nonlinear PDEs with an additive spatio-temporal description of white noise. We study first and second order systems and in doing so, extend several results to the case of second order PDEs. The described approach is based on variational optimization, and performs joint RL-type optimization of the feedback control law and the actuator design over episodes. We demonstrate the efficacy of the proposed approach with several simulated experiments on a variety of SPDEs.

## I. INTRODUCTION

In many complex natural processes, a variable such as temperature or displacement has values that are time varying on a spatial continuum over which the system is defined. These spatio-temporal processes are typically described by Partial Differential Equations (PDEs) and are increasingly prevalent throughout the robotics community. Swarm robotics can be described by Reaction-advection-diffusion PDEs [1]. Robot navigation in crowded environments can be described by Nagumo-like PDEs [2]. Soft robotic limbs can be modelled as damped Euler-Bernoulli systems [3].

Some of the major control-related challenges of these systems include dramatic under-actuation, high system dimensionality, and the design and/or placement of distributed actuators over a continuum of potential locations. These systems often have significant time delay from a control signal, and can have several bifurcations and multi-modal instabilities. In addition, realistic representations of these systems are stochastic. Put together, control of spatio-temporal systems represents many of the largest current-day challenges facing the robotics and automatic control communities.

This paper addresses stochastic optimal control and co-design of Stochastic Partial Differential Equations (SPDEs) through the lens of stochastic optimization. We propose a joint actuator placement and policy network optimization strategy via episodic reinforcement that leverages inherent spatio-temporal stochasticity in the dynamics for optimization. The resulting stochastic gradient descent approach bootstraps

off the widespread success of SGD methods such as ADAM for training Artificial Neural Networks (ANNs).

Among the main goals of this line of research is to establish capabilities for the eventual design and manufacture of soft-body robots. The behavior of such systems follows second order SPDEs. As such, while the proposed method is general to first and second order systems, we focus our mathematical formulation on second order SPDEs.

**Our contributions.** In this work we tackle the coupled challenge of policy optimization and actuator co-design for SPDEs. Our approach is founded on a general principle coming from thermodynamics that also has had success in stochastic optimal control literature [4]

$$\text{Free Energy} \leq \text{Work} - \text{Temperature} \times \text{Entropy} \quad (1)$$

We leverage this principle in order to derive a measure-theoretic loss function that utilizes exponential averaging over importance sampled system trajectories in order to choose network and actuator design parameters that simultaneously minimize state cost and control effort. This work builds off related work [5] by including actuator co-design and establishing needed theory for applicability to second order SPDEs. Specifically, we contribute the following:

- i) A derivation of a joint policy optimization and actuator co-design architecture in Hilbert spaces for SPDEs
- ii) A version of the Girsanov Theorem suitable for second-order SPDEs in a natural form
- iii) A practical set of tools to extend related work to actuator co-design and to handle second-order SPDEs

## II. RELATED WORK

**Control of SPDEs.** There is a growing body of work that seek control of PDEs by immediately reducing them to a set of ODEs [6, 7, 8, 9, 10]. They do not consider stochasticity and typically use standard tools from finite-dimensional control theory. Other approaches apply SPDE control theory, yet either treat linear systems or lack numerical results [11, 12]. We build the proposed framework on [5], wherein the authors create a semi-model-free reinforcement learning framework for policy-based control of SPDEs. This prior work does not consider actuator co-design and lacks a mathematical treatment suitable to second-order SPDEs.

**Actuator Co-Design for PDEs.** Several works have addressed optimal placement of actuators and sensors in the linear

regime, using minimum norm control for the stochastic heat equation [13], using  $H_\infty$  and  $H_2$  objectives for flexible structures [14, 15, 16] and the linearized Ginzburg-Landau equation [17, 18], leveraging symmetry properties in linear PDEs [19], by utilizing Gramians [20, 21], and level set methods based on Gramians that promise scalability [22]. Aside from these methods which are appealing, yet constrained to linear systems, optimal actuator and sensor placement for stabilization of the nonlinear Kuramoto-Sivashinsky equation is demonstrated in [23]. They produce appealing results, however they impose strong simplifying assumptions which limit their dimensionality. Finally, conditions for the existence of optimal actuator and sensor placement for semilinear PDEs are obtained in [24].

### III. PROBLEM FORMULATION

This work proposes actuator co-design optimization for a large class of stochastic spatio-temporal systems represented as SPDEs. We describe these systems as evolving on time separable Hilbert spaces, where they are represented by infinite dimensional vectors and acted on by operators. We address a class of SPDEs that are of *semi-linear* form. Let  $\mathcal{H}$  denote a separable Hilbert space with  $\sigma$ -field  $\mathcal{B}(\mathcal{H})$  and probability space  $(\Omega, \mathcal{F}, \mathbb{P})$  with filtration  $\mathcal{F}_t$ ,  $t \in [0, T]$ . Consider the general semi-linear form of a controlled SPDE on  $\mathcal{H}$  given by

$$dX = (\mathcal{A}X + F(t, X))dt + G(t, X)(\Phi(t, X, \mathbf{x}; \Theta^{(k)})d\mathbf{t} + \frac{1}{\sqrt{\rho}}dW(t)), \quad (2)$$

where  $X(t) \in \mathcal{H}$  is the state of the system which evolves on the Hilbert space  $\mathcal{H}$ , the linear and nonlinear measurable operators  $\mathcal{A} : \mathcal{H} \rightarrow \mathcal{H}$  and  $F(t, X) : \mathbb{R} \times \mathcal{H} \rightarrow \mathcal{H}$  (resp.) are uncontrolled drift terms,  $\Phi(t, X, \mathbf{x}; \Theta^{(k)}) : \mathbb{R} \times \mathcal{H} \times D \rightarrow \mathcal{H}$  is the nonlinear control policy parameterized by  $\Theta^{(k)}$  at the  $k^{\text{th}}$  iteration, where  $D \subset \mathbb{R}^3$  is the domain of the finite spatial region,  $dW(t) : \mathbb{R} \rightarrow \mathcal{H}$  is a spatio-temporal noise process, and  $G(t, X)$  is a nonlinearity that affects both the noise and the control and acts to incorporate each into the field.

The *Hilbert spaces* formulation given in eq. (2) is general in that any semi-linear SPDE can be described in this form by appropriately choosing the  $\mathcal{A}$  and  $F$  operators. In this form, the spatio-temporal noise process  $dW(t)$  is a Hilbert space-valued Wiener process, which is a generalization of the Wiener process in finite dimensions. We include a formal definition of a Wiener process in Hilbert spaces for clarity [25, Section 4.1.1]

**Definition III.1.** A  $\mathcal{H}$ -valued stochastic process  $W(t)$  with probability law  $\mathcal{L}(W(\cdot))$  is called a Wiener process if

- i)  $W(0) = 0$
- ii)  $W$  has continuous trajectories
- iii)  $W$  has independent increments
- iv)  $\mathcal{L}(W(t) - W(s)) = \mathcal{N}(0, t - s)Q$ ,  $t \geq s \geq 0$
- v)  $\mathcal{L}(W(t)) = \mathcal{L}(-W(t))$ ,  $t \geq 0$

**Proposition III.1.** Let  $\{e_i\}_{i=1}^\infty$  be a complete orthonormal system for the Hilbert Space  $\mathcal{H}$ . Let  $Q$  denote the covariance

operator of the Wiener process  $W(t)$ . Note that  $Q$  satisfies  $Qe_i = \lambda_i e_i$ , where  $\lambda_i$  is the eigenvalue of  $Q$  that corresponds to eigenvector  $e_i$ . Then,  $W(t) \in \mathcal{H}$  has the following expansion:

$$W(t) = \sum_{j=1}^{\infty} \sqrt{\lambda_j} \beta_j(t) e_j, \quad (3)$$

where  $\beta_j(t)$  are real valued Brownian motions that are mutually independent on  $(\Omega, \mathcal{F}, \mathbb{P})$ .

This expansion in eq. (3) reveals how the Wiener process acts spatially. There are various forms of the Wiener process with different properties. We refer the interested reader to [25] for a more complete introduction. The proposed approach is derived for a special case of Wiener process called the Cylindrical Wiener process, defined as follows.

**Definition III.2.** A Wiener process  $W(t)$  on  $\mathcal{H}$  is called a Cylindrical Wiener process if the covariance operator  $Q$  is the identity operator  $I$ .

Note that for the Cylindrical Wiener process, the sum in eq. (3) is unbounded in  $\mathcal{H}$  since  $\lambda_j = 1$ ,  $\forall j = 1, 2, \dots$ . This makes the Cylindrical Wiener process a challenging Wiener process to handle since it acts spatially *everywhere* with *equal* magnitude, in contrast to Wiener processes with covariance operators that are of trace class (i.e. wherein the expansion eq. (3) is finite). This type of spatio-temporal noise requires the assumption that the operators  $\mathcal{A}$ ,  $F(t, X)$ , and  $G(t, X)$  satisfy properly formulated conditions given in [25, Hypothesis 7.2] to guarantee the existence and uniqueness of the  $\mathcal{F}_t$ -adapted weak solution  $X(t)$ ,  $t \geq 0$ .

The nonlinear policy  $\Phi(t, X, \mathbf{x}; \Theta^{(k)})$  is a potentially time-varying policy that has explicit state dependence. Nonlinear, explicit state dependence allows for a feedback policy that can extract pertinent information from the state for control. In this work the nonlinear policy utilizes an ANN. Embedded in this function is also a spatial dependence. The dependence on  $\mathbf{x}$  describes how the actuator is placed in the spatial domain. This approach encompasses cases where terms that parametrize how the actuators are shaped or sized are included in the nonlinear policy.

Many complex spatio-temporal systems are given by partial differential equations of second order in time. One such system is the simply supported stochastic Euler-Bernoulli equation with Kelvin-Voigt and viscous damping, which can describe the motion of a soft robotic limb. Formally, this is given by

$$\begin{aligned} \partial_{tt}y + \partial_{xx}(\partial_{xx}y + C_d \partial_{xxt}y) + \mu \partial_t y &= \Phi + \frac{1}{\sqrt{\rho}} \partial_t W(t, x) & (4) \\ y(t, 0) = y(t, a) = 0, \quad y(0, x) = y_0, \quad \partial_t y(0, x) &= v_0 \\ \partial_{xx}y(t, 0) + C_d \partial_{xxt}y(t, 0) = 0, \quad \partial_{xx}y(t, a) + C_d \partial_{xxt}y(t, a) &= 0 \end{aligned}$$

where the spatial region is one-dimensional,  $y(x, t) = y : \mathbb{R}^n \times \mathbb{R} \rightarrow \mathbb{R}$  represents the vertical displacement of the beam, and all functional dependencies of the nonlinear policy  $\Phi$  have been dropped since it has a different form in the PDE perspective. With the change of variables  $v := \partial_t y$ , this system has the

typical second order matrix form

$$\begin{aligned} \partial_t \begin{bmatrix} y \\ v \end{bmatrix} &= \begin{bmatrix} 0 & 1 \\ -A_0 & -C_d A_0 - \mu \end{bmatrix} \begin{bmatrix} y \\ v \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} \Phi \\ &+ \begin{bmatrix} 0 \\ \frac{1}{\sqrt{\rho}} \end{bmatrix} \partial_t W(t, x) \end{aligned} \quad (5)$$

where  $A_0 = \partial_{xxx}$  without boundary conditions. Now, we lift this PDE into infinite dimensional Hilbert spaces. Define  $Y \in \mathcal{H}$  as the Hilbert space analog of  $y(x, t)$ ,  $V \in \mathcal{H}$  as the Hilbert space analog of  $v(x, t)$ , and a variable  $Z$  on the direct product Hilbert space  $\mathcal{H}^2 := \mathcal{H} \times \mathcal{H}$ . Note that  $Z$  is a Hilbert space analog of a variable  $z(x, t) = [y(x, t) \ v(x, t)]^\top \in \mathbb{R}^2$ . In Hilbert spaces,  $A_0$  becomes an operator acting on  $\mathcal{H}$  and 1 gets replaced by the identity operator  $I$  acting on  $\mathcal{H}$ . Rewriting eq. (5) in Hilbert space semi-linear form yields

$$dZ = AZ dt + G\left(\Phi(t, Z, \mathbf{x}; \Theta^{(k)})\right) dt + \frac{1}{\sqrt{\rho}} dW(t) \quad (6)$$

where  $A: \mathcal{H}^2 \rightarrow \mathcal{H}^2$  is the linear operator  $A = [0 \ I; -A_0 \ -C_d A_0 - \mu I]$ ,  $G: \mathcal{H} \rightarrow \mathcal{H}^2$  is operator representing how control and spatio-temporal noise enter the system  $G = [0; I]$ , and  $dW(t)$  is a Cylindrical Wiener process on  $\mathcal{H}$ . Note that the Hilbert space variables  $Y$ ,  $V$ , and  $Z$  no longer have spatial dependence as the Hilbert space vectors capture the spatial continuum over which the problem is defined.

#### IV. A GIRSANOV THEOREM FOR SECOND ORDER SPDES

In order to derive a measure theoretic view of variational optimization, we require a change of measures coming from an appropriate Girsanov theorem for second order Hilbert space valued systems. We present the second order version of the Girsanov theorem here as an extension of existing formulations, where all norms and inner products are taken with respect to the Hilbert space  $\mathcal{H}$ .

**Theorem IV.1** (Girsanov). *Let  $\Omega$  be a sample space with a  $\sigma$ -algebra  $\mathcal{F}$ . Consider the following  $\mathcal{H}^2$ -valued nonlinear stochastic processes*

$$dZ = (\mathcal{A}Z + F(t, Z)) dt + \frac{1}{\sqrt{\rho}} G(t, Z) dW(t), \quad (7)$$

$$d\tilde{Z} = (\mathcal{A}\tilde{Z} + F(t, \tilde{Z})) dt + G(t, \tilde{Z}) \left( \tilde{B}(t, \tilde{Z}) dt + \frac{1}{\sqrt{\rho}} dW(t) \right), \quad (8)$$

where  $Z(0) = \tilde{Z}(0) = z_0$  and  $W \in \mathcal{H}$  is a Cylindrical Wiener process with respect to measure  $\mathbb{P}$ . Moreover, let  $\Gamma$  be a set of continuous-time, infinite-dimensional trajectories in the time interval  $[0, T]$ . Define the probability law of  $Z$  over trajectories  $\Gamma$  as  $\mathcal{L}(\Gamma) := \mathbb{P}(\omega \in \Omega | Z(\cdot, \omega) \in \Gamma)$ . Similarly, define the law of  $\tilde{Z}$  as  $\tilde{\mathcal{L}}(\Gamma) := \mathbb{P}(\omega \in \Omega | \tilde{Z}(\cdot, \omega) \in \Gamma)$ . Then

$$\begin{aligned} \tilde{\mathcal{L}}(\Gamma) &= \mathbb{E}_{\mathbb{P}} \left[ \exp \left( \int_0^T \langle \psi(s), dW(s) \rangle \right. \right. \\ &\quad \left. \left. - \frac{1}{2} \int_0^T \|\psi(s)\|^2 ds \right) \Big| X(\cdot) \in \Gamma \right], \end{aligned} \quad (9)$$

where we have defined

$$\psi(t) := \sqrt{\rho} \tilde{B}(t, Z(t)) \in \mathcal{H}, \quad (10)$$

and assumed  $\mathbb{E}_{\mathbb{P}} [e^{\frac{1}{2} \int_0^T \|\psi(t)\|^2 dt}] < +\infty$ .

*Proof:* Define the process

$$\hat{W}(t) := W(t) - \int_0^t \psi(s) ds. \quad (11)$$

Under the above assumption,  $\hat{W}$  is a Cylindrical Wiener process with respect to a measure  $\mathbb{Q}$  defined by

$$\begin{aligned} d\mathbb{Q}(\omega) &= \exp \left( \int_0^T \langle \psi(s), dW(s) \rangle - \frac{1}{2} \int_0^T \|\psi(s)\|^2 ds \right) d\mathbb{P} \\ &= \exp \left( \int_0^T \langle \psi(s), d\hat{W}(s) \rangle + \frac{1}{2} \int_0^T \|\psi(s)\|^2 ds \right) d\mathbb{P}. \end{aligned} \quad (12)$$

The proof for this intermediate result can be found in [26, Theorem 10.14]. Now, using eq. (11), eq. (7) is rewritten as

$$\begin{aligned} dZ &= (\mathcal{A}Z + F(t, Z)) dt + \frac{1}{\sqrt{\rho}} G(t, Z) dW(t) \\ &= (\mathcal{A}Z + F(t, Z)) dt + G(t, Z) \left( B(t, Z) dt + \frac{1}{\sqrt{\rho}} d\hat{W}(t) \right) \end{aligned} \quad (13)$$

Notice that the SPDE in eq. (14) has the same form as eq. (8). Therefore, under the introduced measure  $\mathbb{Q}$  and noise profile  $\hat{W}$ ,  $Z(\cdot, \omega)$  becomes equivalent to  $\tilde{Z}(\cdot, \omega)$ . Conversely, under measure  $\mathbb{P}$ , eq. (13) (or eq. (14)) behaves as the original system in eq. (7). In other words, eq. (7) and eq. (14) describe the same system on  $(\Omega, \mathcal{F}, \mathbb{P})$ . From the uniqueness of solutions and the aforementioned reasoning, one has

$$\mathbb{P}(\{\tilde{Z} \in \Gamma\}) = \mathbb{Q}(\{Z \in \Gamma\}).$$

The result follows from eq. (12).  $\blacksquare$

It follows that the *Randon-Nikodym* (RN) derivative between measures  $\mathcal{L}(\cdot)$  and  $\tilde{\mathcal{L}}(\cdot)$  of the different dynamical systems defined in eqs. (7) and (8), is given by

$$\frac{d\mathcal{L}}{d\tilde{\mathcal{L}}} = \exp \left( - \int_0^T \langle \psi(s), dW(s) \rangle - \frac{1}{2} \int_0^T \|\psi(s)\|^2 ds \right). \quad (15)$$

Applying this to both the general semi-linear system eq. (2) and the second order semi-linear form of the Euler-Bernoulli system eq. (6) yields  $\psi := \sqrt{\rho} \Phi(t, Z, \mathbf{x}; \Theta^{(k)})$  and a Radon-Nikodym derivative as

$$\begin{aligned} \frac{d\mathcal{L}}{d\tilde{\mathcal{L}}} &= \exp \left( - \sqrt{\rho} \int_0^T \langle \Phi(t, Z, \mathbf{x}; \Theta^{(k)}), dW(s) \rangle \right. \\ &\quad \left. - \frac{\rho}{2} \int_0^T \|\Phi(t, Z, \mathbf{x}; \Theta^{(k)})\|^2 ds \right). \end{aligned} \quad (16)$$

Throughout the rest of the manuscript, we refer to the terms in the Radon-Nikodym derivative many times. For convenience, they will be assigned functions  $\mathcal{N}(\Theta, \mathbf{x}) := \int_0^T \langle \Phi(t, Z, \mathbf{x}; \Theta^{(k)}), dW(s) \rangle$ , and  $\mathcal{P}(\Theta, \mathbf{x}) := \int_0^T \|\Phi(t, Z, \mathbf{x}; \Theta^{(k)})\|^2 ds$ .

## V. SPATIO-TEMPORAL STOCHASTIC OPTIMIZATION

The proposed measure theoretic framework is based on eq. (1) in the following form [27, 28]

$$-\frac{1}{\rho} \log \mathbb{E}_{\mathcal{L}} \left[ \exp(-\rho J) \right] = \min_{\Theta, \mathbf{x}} \left[ \mathbb{E}_{\tilde{\mathcal{L}}} (J) + \frac{1}{\rho} D_{KL}(\tilde{\mathcal{L}} \parallel \mathcal{L}) \right], \quad (17)$$

where  $J = J(X)$  is an arbitrary state cost function. Relating eq. (17) to eq. (1), the metaphorical work and entropy describe a metaphorical energy landscape for which there is a minimizing measure. Sampling from this measure would simultaneously minimize state cost and the  $KL$ -divergence term, which is interpreted as control effort. The measure that optimizes eq. (17) is the so-called Gibbs measure

$$d\mathcal{L}^* = \frac{\exp(-\rho J) d\mathcal{L}}{\mathbb{E}_{\mathcal{L}} [\exp(-\rho J)]}. \quad (18)$$

It is not known how to sample directly from the Gibbs measure in eq. (18). Instead, variational optimization methods seek to iteratively minimize the controlled distribution's distance to the Gibbs measure [28, 29, 30]. Define the control policy and actuator co-design problem as

$$\Theta^* = \operatorname{argmin}_{\Theta} D_{KL}(\mathcal{L}^* \parallel \tilde{\mathcal{L}}) \quad (19a)$$

$$\mathbf{x}^* = \operatorname{argmin}_{\mathbf{x}} D_{KL}(\mathcal{L}^* \parallel \tilde{\mathcal{L}}) \quad (19b)$$

Throughout experiments, the authors found that a joint optimization problem dramatically outperforms the split problem in eq. (19). To make this clear, define a new variable  $\hat{\Theta} := [\Theta; \mathbf{x}]^\top$ , and with it the new joint variational optimization as

$$\hat{\Theta}^* = \operatorname{argmin}_{\hat{\Theta}} D_{KL}(\mathcal{L}^* \parallel \tilde{\mathcal{L}}). \quad (20)$$

Expanding the  $KL$  divergence and applying the chain rule yields

$$\hat{\Theta}^* = \operatorname{argmin}_{\hat{\Theta}} \left[ \int \log \left( \frac{d\mathcal{L}^*}{d\mathcal{L}} \frac{d\mathcal{L}}{d\tilde{\mathcal{L}}} \right) d\mathcal{L}^* \right], \quad (21)$$

which is equivalent to minimizing

$$\hat{\Theta}^* = \operatorname{argmin}_{\hat{\Theta}} \left[ \int \log \left( \frac{d\mathcal{L}}{d\tilde{\mathcal{L}}} \right) d\mathcal{L}^* \right]. \quad (22)$$

Performing importance sampling yields

$$\hat{\Theta}^* = \operatorname{argmin}_{\hat{\Theta}} \left[ \int \log \left( \frac{d\mathcal{L}}{d\tilde{\mathcal{L}}} \right) \frac{d\mathcal{L}^*}{d\mathcal{L}} \frac{d\mathcal{L}}{d\tilde{\mathcal{L}}} d\tilde{\mathcal{L}} \right]. \quad (23)$$

The proposed iterative approach performs episodic reinforcement with respect to a loss function in order to optimize eq. (23). Define the loss function as

$$L(\hat{\Theta}) := \mathbb{E}_{\tilde{\mathcal{L}}} \left[ \log \left( \frac{d\mathcal{L}}{d\tilde{\mathcal{L}}} \right) \frac{d\mathcal{L}^*}{d\mathcal{L}} \frac{d\mathcal{L}}{d\tilde{\mathcal{L}}} \right] \quad (24)$$

Plugging eqs. (15) and (18) into eq. (24) yields

$$L(\hat{\Theta}^{(k)}) = \mathbb{E}_{\tilde{\mathcal{L}}} \left[ \frac{\exp(-\rho \tilde{J})}{\mathbb{E}_{\tilde{\mathcal{L}}} [\exp(-\rho \tilde{J})]} \left( -\sqrt{\rho} \mathcal{N}(\hat{\Theta}^{(k)}) - \frac{\rho}{2} \mathcal{P}(\hat{\Theta}^{(k)}) \right) \right], \quad (25)$$

where  $\tilde{J} = \tilde{J}(Z_0^T, \hat{\Theta}^{(k)})$  is defined as

$$\tilde{J}(Z_0^T, \hat{\Theta}^{(k)}) := J(Z_0^T) + \frac{1}{\sqrt{\rho}} \mathcal{N}(\hat{\Theta}^{(k)}) + \frac{1}{2} \mathcal{P}(\hat{\Theta}^{(k)}), \quad (26)$$

and  $J(Z_0^T)$  is a state cost evaluated over the state trajectory  $Z_0^T$ . For reaching tasks,  $J(Z_0^T)$  is typically a weighted quadratic penalization of the 1-norm distance to the goal state.

This loss function compares sampled trajectories by evaluating them on the exponentiated  $\tilde{J}$  performance metric. The importance sampling terms  $\mathcal{N}$  and  $\mathcal{P}$ , which appear in  $\tilde{J}$  add a quadratic control penalization term and a mixed control noise term. In the Loss function, they serve as weights for the exponentiated cost trajectories. For convenience, we denote the exponentiated cost term as  $\mathcal{E} := \mathbb{E}_{\tilde{\mathcal{L}}} [\exp(-\rho \tilde{J})]^{-1} \exp(-\rho \tilde{J})$ .

Recall, that the nonlinear policy  $\Phi$  is a functional mapping into Hilbert space  $\mathcal{H}$ . This is kept general for derivation purposes, however it implies that the nonlinear policy controls each element of an infinite vector  $Z \in \mathcal{H}$ . A more realistic, but less general representation refines the policy as

$$\Phi(t, Z, \mathbf{x}; \Theta^{(k)}) = \mathbf{m}(\mathbf{x})^\top \varphi(Z; \Theta^{(k)}), \quad (27)$$

where  $\mathbf{m}(\mathbf{x}) : D^N \rightarrow \mathbb{R}^N \times \mathcal{H}$  represents the effect of the actuation from  $N$  actuators on the infinite-dimensional field. Typically this is either a Gaussian-like exponential with mean centered at the actuator locations or an indicator function.

In eq. (27),  $\varphi(X; \Theta^{(k)}) : \mathcal{H} \rightarrow \mathbb{R}^N$  is a policy network with  $N$  control outputs representing  $N$  distributed (or boundary) actuators. Note that as desired, the tensor contraction given on the right hand side of eq. (27) produces a vector in  $\mathcal{H}$ . Splitting the actuation function from the control signal is also desired because we ultimately wish to use a finite input, finite output policy network for the function  $\varphi(X; \Theta^{(k)})$ . The inner product terms become

$$\mathcal{N}(\hat{\Theta}^{(k)}) = \int_0^T \langle \mathbf{m}(\mathbf{x})^\top \varphi(Z; \Theta^{(k)}), dW(s) \rangle \quad (28)$$

$$\begin{aligned} \mathcal{P}(\hat{\Theta}^{(k)}) &= \int_0^T \|\mathbf{m}(\mathbf{x})^\top \varphi(Z; \Theta^{(k)})\|^2 ds \\ &= \int_0^T \langle \varphi(Z; \Theta^{(k)}), \mathbf{M}(\mathbf{x}) \varphi(Z; \Theta^{(k)}) \rangle ds \end{aligned} \quad (29)$$

where  $\mathbf{M}(\mathbf{x}) = \mathbf{m}(\mathbf{x})\mathbf{m}(\mathbf{x})^\top$ .

## VI. APPROXIMATE DISCRETE OPTIMIZATION

Performing spatio-temporal stochastic optimization in Hilbert spaces as described above maintains generality of the resulting loss function. In addition, it allows us to make use of the spatio-temporal noise process. Finally, we can apply the spatial inner product integration scheme described in [5]. Despite this, computations on a digital computer require spatial discretization.

As a result, the portion of this optimization procedure dedicated to actuator co-design becomes a discrete optimization problem.

To see this more clearly, consider the 1-dimensional (1D) spatial continuum  $D = [0, 1]$  discretized into a 10 point 1D grid. Lets assume that an actuator is chosen to be placed at  $x = 0.25$ . Even though the actuation function  $m(x)$  may be Gaussian-like function, the majority of the actuation will be felt in between two grid points, namely 0.2 and 0.3. This problem is even more severe if the actuation function  $m(x)$  is the indicator function, as there will be no actuation exerted on the field irrespective of the control signal magnitude. Denote the number of spatial discretization points as  $J$  and the discretized problem domain grid as  $\hat{D}$  composed of  $J^3$  elements. The optimization problem becomes

$$\begin{aligned} \min_{\Theta, \mathbf{x}} L(\Theta, \mathbf{x}) \\ \text{subject to } \mathbf{x} \in \hat{D} \end{aligned} \quad (30)$$

This formulation is an accurate representation, yet limits gradient flow from the loss function back to the actuator design parameters. In order to maintain these gradients, we approximate eq. (30) as follows. Define a one-to-one map  $S: \hat{D} \rightarrow \mathbb{Z}_+$ , where  $\mathbb{Z}_+$  denotes positive integers. Applying the forward and inverse mapping produces a gradient-based parameter update of the form

$$\Theta^{(k+1)} = \Theta^{(k)} - \gamma_{\Theta} \nabla_{\Theta} L(\Theta^{(k)}, \mathbf{x}^{(k)}) \quad (31)$$

$$\mathbf{x}^{(k+1)} = S^{-1} \left( R \left( S(\mathbf{x}^{(k)} - \gamma_{\mathbf{x}} \nabla_{\mathbf{x}} L(\Theta^{(k)}, \mathbf{x}^{(k)})) \right) \right) \quad (32)$$

where  $R(\cdot)$  simply rounds to the nearest integer,  $\gamma_{\Theta}$  is the learn rate for the ANN parameters,  $\gamma_{\mathbf{x}}$  is the learn rate for the actuator design parameters,  $\nabla_{\Theta}$  denotes the gradient with respect to  $\Theta$ , and  $\nabla_{\mathbf{x}}$  denotes the gradient with respect to  $\mathbf{x}$ . This approach allows us to leverage well-known backprop-based algorithms such as ADA-Grad [31] and ADAM [32].

## VII. ALGORITHM AND NETWORK ARCHITECTURE

As discussed previously, implementation of the above framework requires spatial and temporal discretization of the SPDEs discussed in section III. With this in mind, we choose an ANN for our nonlinear policy  $\varphi(Z; \Theta^{(k)})$ . In this work we exclusively use Feed-forward Neural Networks (FNNs) for all of our experiments, and use physics-based models of each SPDEs to generate training data. Given that the proposed framework is semi-model-free, real system data can seamlessly replace the physics-based model as described in [5]. We only need prior knowledge of the flavor of noise and the actuator dynamics (i.e.  $\mathbf{m}(\mathbf{x})$ ). Also described is a sparse method for spatial integration that we apply here to maintain computational and memory efficiency.

The resulting algorithm, which we call the Actuator Design and Policy Learning (ADPL) algorithm is shown in algorithm 1. The inputs are time horizon ( $T$ ), number of iterations ( $K$ ), number of rollouts ( $R$ ), initial state ( $Z_0$ ), number of actuators ( $N$ ), noise variance ( $\rho$ ), time discretization ( $\Delta t$ ), actuator variance ( $\sigma_{\mu}$ ), initial network parameters ( $\Theta^{(0)}$ ), initial actuator

---

## Algorithm 1 Actuator Design and Policy Learning

---

```

1: Function:  $\Theta^* = \text{OptimizePolicyNetwork}(T, K, R, Z_0, N, \rho, \Delta t, \mu, \sigma_{\mu}, \Theta^{(0)}, \mathbf{x}^{(0)}, \gamma_{\Theta}, \gamma_{\mathbf{x}})$ 
2: for  $k = 1$  to  $K$  do
3:   Compute  $\mathbf{m}(\mathbf{x}), M(\mathbf{x}) \forall \mathbf{y} \in D$ 
4:   for  $r = 1$  to  $R$  do
5:     for  $t = 1$  to  $T$  do
6:        $dW_t \leftarrow \text{SampleNoise}()$ 
7:        $Z_t \leftarrow \text{Propagate}(Z_{t-1}, \Theta^{(k)}, dW_t)$  via eq. (6)
8:        $u_t \leftarrow \text{SparseForwardPass}(\Theta^{(k)}, Z_t)$ 
9:     end for
10:   end for
11:    $J \leftarrow \text{StateCost}(Z_0^T)$ 
12:    $N \leftarrow \mathcal{N}(u_0^T, dW, \mathbf{m}(\mathbf{x}))$  via eq. (28)
13:    $P \leftarrow \mathcal{P}(u_0^T, M(\mathbf{x}))$  via eq. (29)
14:    $E \leftarrow \mathcal{E}(J, N, P)$  as in eq. (25)
15:    $L \leftarrow \text{ComputeLoss}(P, N, E)$  via eq. (26)
16:   Compute  $\nabla_{\Theta}(L)$  via backprop
17:   Compute  $\nabla_{\mathbf{x}}(L)$  via backprop
18:    $\Theta^{(k+1)} \leftarrow \text{GradientStep}(L, \gamma_{\Theta})$  via eq. (31)
19:    $\mathbf{x}^{(k)} \leftarrow \text{GradientStep}(L, \gamma_{\mathbf{x}})$  via eq. (32)
20: end for

```

---

locations ( $\mathbf{x}^{(0)}$ ), policy learn rate ( $\gamma_{\Theta}$ ), and actuator location learn rate ( $\gamma_{\mathbf{x}}$ ). For more information on *SampleNoise()*, refer to [33, Chapter 10].

In addition to obtaining gradients from any variant of Stochastic Gradient Descent (SGD) via backprop paths, *GradientStep* for actuator parameters also adds the current gradient to the previous gradient if the actuator location has not moved by the current gradient step. This is useful because when the actuators are near their optimal value, the gradients become very small, preventing the actuator locations from reaching optimal values. This could also be achieved by heuristically changing the learning rate  $\gamma_{\mathbf{x}}$ . Also, a quadratic cost term was added to  $L$  in eqs. (31) and (32) in order to penalize actuators from leaving the spatial domain  $\hat{D}$ , with a large coefficient to ensure this condition was rarely, if ever violated.

Note that there are separate learning rates for the policy and the actuators. This is because in practice the authors found that the optimization landscape is typically much more shallow for the actuator design than for the policy parameters. For most of the experiments, the actuator placement learn rate  $\gamma_{\mathbf{x}}$  was about 30 times larger than the policy network learn rate  $\gamma_{\Theta}$ .

## VIII. SIMULATION RESULTS AND DISCUSSION

We applied our approach to four simulated SPDE experiments to simultaneously place actuators and optimize a policy network. Each experiment used less than 32 GB RAM, and was run on a desktop computer with a Intel Xeon 12-core CPU with a NVIDIA GeForce GTX 980 GPU. Our code was written to operate inside a Tensorflow graph [34] to leverage rapid static graph computation, as well as sparse linear algebra operations used by *SparseForwardPass* [5]. The first two experiments involved a reaching task, where the SPDEs is initialized at a

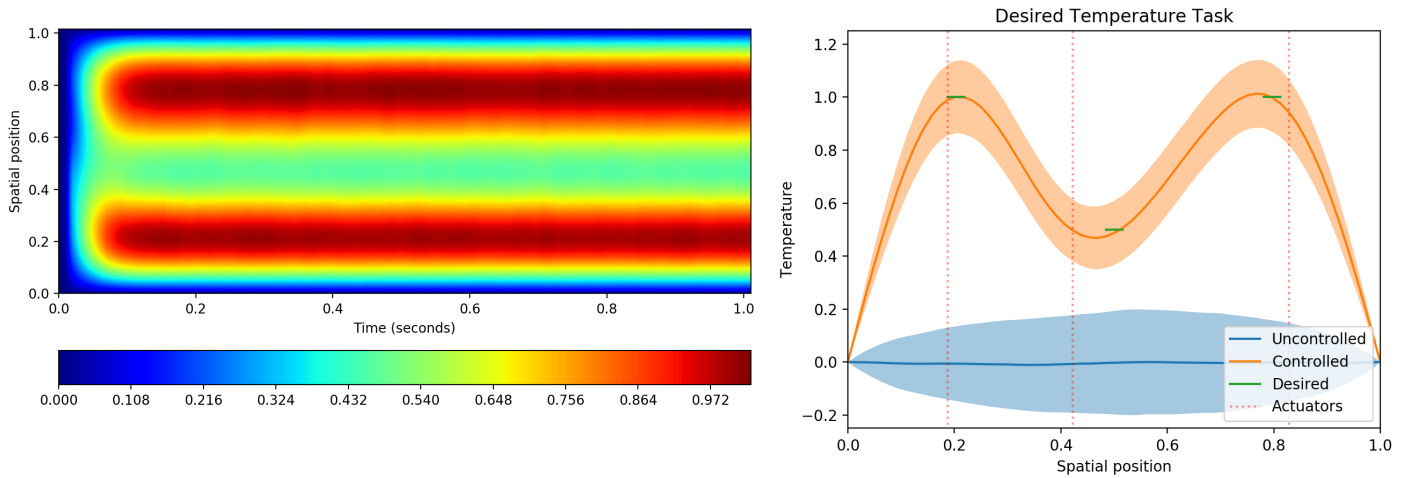


Fig. 1: Heat Equation Temperature Reaching Task. (left) controlled contour plot where color represents temperature, (right) final time snapshot comparing to the uncontrolled system. Mean trajectories are represented with a solid line, while a  $2\sigma$  standard deviation is represented with a shaded region.

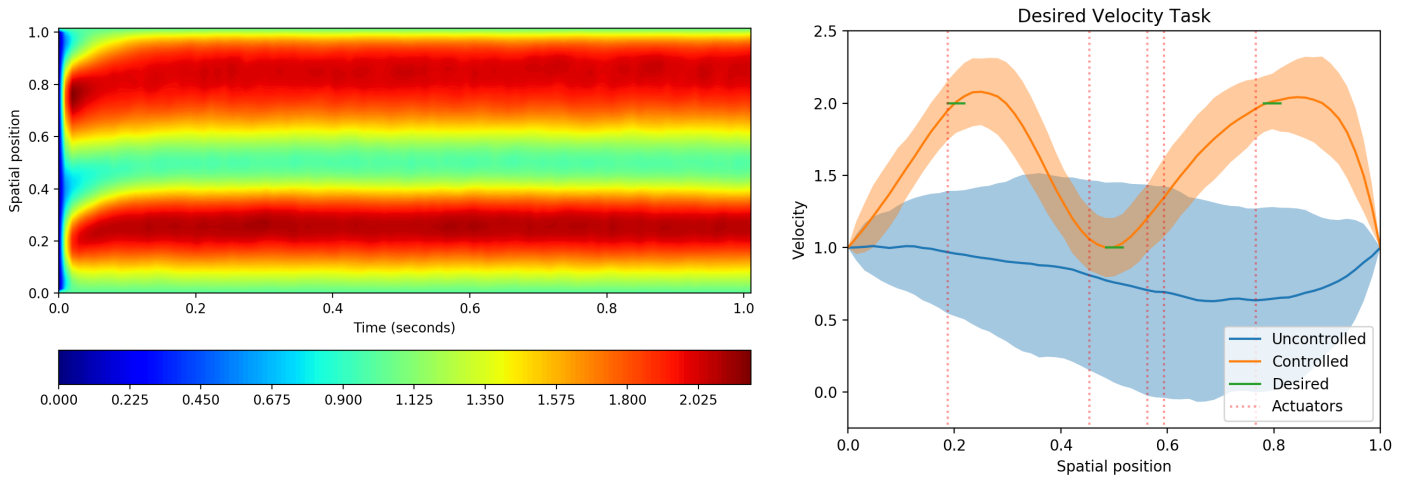


Fig. 2: Burgers Velocity Reaching Task. (left) controlled contour plot where color represents velocity, (right) final time snapshot comparing to the uncontrolled system. Mean trajectories are represented with a solid line, while a  $2\sigma$  standard deviation is represented with a shaded region.

zero initial condition over the spatial region, and must reach certain values at pre-specified regions of the spatial domain. The last two experiments involved a suppression task, where some non-zero initial condition must be suppressed on desired regions.

The data that was used for training was generated by a spatial central difference, semi-implicit time discretized version of each SPDE. These schemes are described in detail in [33, Chapter 3 & 10]. Each experiment had all actuators initialized to be sampled from a random distribution on  $[0.4a, 0.6a]$ , where  $a$  denotes the spatial size. For 3500 iterations of our algorithm, run times for the most complicated system—the Euler-Bernoulli equation—were about 15 hours. Details of the experiments and videos of the controlled systems can be found in the provided links<sup>1</sup>. We encourage the interested reader to contact the authors

for code.

The first experiment was a temperature reaching task on the 1D Heat equation with homogeneous,  $T = 0$  Dirichlet boundary conditions, and is depicted in fig. 1. The task was to raise the temperature at regions specified in green to specified values depicted in the figure. The algorithm was run for 3000 iterations.

The next experiment was a velocity reaching task on the Burgers equation with non-homogenous Dirichlet boundary conditions, and is depicted in fig. 2. The Burgers equation has a nonlinear advection term, which produces an apparent rightward motion. The algorithm was run for 3500 iterations, and was able to take advantage of the advection for actuator placement in order to solve the task with lower control effort.

The heat equation is a pure diffusion SPDE, while the Burgers equation shares the diffusion term with the Heat equation with an added advection term. The results of the

<sup>1</sup>Supplement: [tinyurl.com/yc7fq3lc](https://tinyurl.com/yc7fq3lc) | Video: <https://youtu.be/pqryLc4wCuU>




Fig. 3: Nagumo Suppression Task. (left) controlled contour plot where color represents voltage, (right) final time snapshot comparing to the uncontrolled system. Mean trajectories are represented with a solid line, while a  $2\sigma$  standard deviation is represented with a shaded region.




Fig. 4: Nagumo Suppression Task Comparison Plots. (left) controlled state cost plot, where solid lines denote mean, and shaded regions denote a  $2\sigma$  standard deviation, (center) Control signal comparison plot, where lines represent mean behavior, and (right) Final time snapshot comparing the actuators placed by our approach and actuators placed by a human expert with policy optimization by IDVRL.




Fig. 5: Euler-Bernoulli Suppression Task. (left) Uncontrolled contour plot (right) Controlled contour plot. In both plots, color represents deflection on top, and deflection velocity on bottom.

Heat and Burgers experiments show actuator locations that take advantage of the natural behavior of each SPDEs. In the case of the Heat equation, actuators are nearby the desired regions such that the temperature profile can reach a flat peak of the diffusion at the desired profile. In the case of the Burgers equation, the advection pushes towards the right end of the space, thus forming a wave front that develops at the right end, but leaves the left end dominated by the diffusion term. This is again reflected in the placement of actuators. The first actuator is nearby the desired region just as the actuators in the Heat SPDE, while two of the actuators between the center and the right region are located to be able to control the amplitude and shape of the developing wave front so as to produce a flat peak that aligns with the desired region at the desired velocity. The central desired region is flanked on both sides by actuators that are nearly equidistant, in order to produce another desired flat velocity region at this location.

The third experiment was a voltage suppression task on the Nagumo equation with homogeneous Neumann boundary conditions, and is depicted in fig. 3. The task was to suppress an initial voltage on the left end, that without intervention propagates toward the right end, as shown by the uncontrolled trajectories. The Nagumo equation is composed of a diffusive term and a 3rd-order nonlinearity, making this equation the most challenging from a nonlinear control perspective. Despite this, our approach was able to simultaneously place actuators and provide control such that the task was solved. The algorithm was run for 2000 iterations, and demonstrates actuator placement optimization that takes advantage of the natural system behavior. This task was also the most challenging due to the significantly longer planning horizon of 3.5 seconds, as compared to the 1.0 second planning horizon of all the other experiments.

In order to validate our proposed approach, we compared the actuator locations that the algorithm found after optimization to the actuator locations that were hand placed by a human expert for the simulated experiments conducted in related work [5]. To have a valid comparison, we ran the IDVRL algorithm for both sets of actuator locations. Figure 4 reports these results. The left figure shows that the state costs for each is almost identical. Note that the scale here is  $10^{-4}$ . The center figure shows the control signals for each actuator, for each method, and demonstrates that for almost identical state cost values, the control effort for each actuator with our approach is lower on average. The calculated average control signal magnitudes for IDVRL are 3.3 times higher than our method. The third plot shows the voltage profile at the final time. We hypothesize that the lower control effort is due to the control over the shape of the spatially propagating signal, enabling it to have a smoother transition into the desired region. While the penalty of this actuator placement is a slightly higher variance on the desired region, the choice appears correct given the result.

The final task was an oscillation suppression task on the Euler-Bernoulli equation with Kelvin-Voigt damping given in eq. (4), and is depicted in fig. 5. As shown, the initial condition prescribes spatial oscillations, that then oscillate

temporally. The second-order nature of the system creates offset and opposite oscillations in the velocity profile, that in turn produce offset and opposite oscillations in the position profile. Without interference, the oscillations proceed over the entire time window. As shown on the right, our approach successfully suppresses these oscillations, which die out quickly under the given control policy. In this experiment, the actuators did not leave the initialized actuator placement region  $[0.4, 0.6]$  prescribed for all experiments.

The Euler-Bernoulli oscillation suppression task is in fact very challenging and complex. Producing a control signal at an actuator location that is in phase with the velocity oscillations will amplify the oscillations, leading to a divergence. The actuator location and control signal from the policy network must work in concert to produce a control signal out of phase with the velocity that matches its frequency, which is time varying due to control, as shown on the right side of fig. 5.

Each of the above experiments has its challenges and in most cases the spatio-temporal problem space produces a joint policy optimization and actuator co-design problem that is littered with local minima. These experiments demonstrate that the proposed approach can jointly optimize a policy network and actuator design. These results and the overall performance of the algorithm are indicative that this approach may enable actuator design on problem spaces where a human has little to no prior knowledge to rely on in attempting to solve the problem by hand.

## IX. CONCLUSION AND FUTURE DIRECTIONS

This work presents a framework for joint policy optimization and actuator co-design. We contribute a novel mathematical tool required for a measure theoretic treatment of second order SPDEs, namely the change of measure and associated proof. We also provide a mechanism for handling discrete parameter optimization with a gradient-based approach. We demonstrate the resulting algorithm on four different SPDEs, each with their own challenges and complexities. The last of which is the Euler-Bernoulli SPDE which connects back with our goal of establishing capabilities for the further development of soft-body robotics.

The presented approach is a new way of performing optimization and can lead to many applications in soft robotics, soft materials, morphing, and continuum mechanism. The results are encouraging to the authors. We plan to scale this approach on systems with higher spatial dimensionality (i.e. 2D and 3D), as well as investigate more complex forms of the above systems that are relevant to a future of soft robotics.

## ACKNOWLEDGEMENTS

This work was supported by Amazon AWS, NSF CMMI #1662523, and by NASA Langley. Ethan N. Evans was supported by the SMART scholarship and George I. Boutsellis was supported by the A. S. Onasis Foundation.



## REFERENCES

- [1] Karthik Elamvazhuthi, Hendrik Kuiper, and Spring Berman. Pde-based optimization for stochastic mapping and coverage strategies using robotic ensembles. *Automatica*, 95:356–367, 2018.
- [2] Eugene Aidman, Vladimir Ivancevic, and Andrew Jennings. A coupled reaction-diffusion field model for perception-action cycle with applications to robot navigation. *International Journal of Intelligent Defence Support Systems*, 1(2):93–115, 2008.
- [3] Yoel Shapiro, Kosa Gabor, and Alon Wolf. Modeling a hyperflexible planar bending actuator as an inextensible euler-bernoulli beam for use in flexible robots. *Soft Robotics*, 2(2):71–79, 2015.
- [4] E.A Theodorou and E. Todorov. Relative entropy and free energy dualities: Connections to path integral and kl control. In *the Proceedings of IEEE Conference on Decision and Control*, pages 1466–1473, Dec 2012.
- [5] Ethan N. Evans, Marcus A. Pereira, George I. Boutselis, and Evangelos A. Theodorou. Variational optimization based reinforcement learning for infinite dimensional stochastic systems. In *Conference on Robot Learning*, 2019.
- [6] Katharina Bieker, Sebastian Peitz, Steven L Brunton, J Nathan Kutz, and Michael Dellnitz. Deep model predictive control with online learning for complex physical systems. *arXiv preprint arXiv:1905.10094*, 2019.
- [7] Aditya G Nair, Chi-An Yeh, Eurika Kaiser, Bernd R Noack, Steven L Brunton, and Kunihiko Taira. Cluster-based feedback control of turbulent post-stall separated flows. *Journal of Fluid Mechanics*, 875:345–375, 2019.
- [8] Arvind T Mohan and Datta V Gaitonde. A deep learning based approach to reduced order modeling for turbulent flow control using lstm neural networks. *arXiv preprint arXiv:1804.09269*, 2018.
- [9] Jeremy Morton, Antony Jameson, Mykel J Kochenderfer, and Freddie Witherden. Deep dynamical modeling and control of unsteady fluid flows. In *Advances in Neural Information Processing Systems*, pages 9258–9268, 2018.
- [10] Jean Rabault, Miroslav Kuchta, Atle Jensen, Ulysse Réglade, and Nicolas Cerardi. Artificial neural networks trained through deep reinforcement learning discover control strategies for active flow control. *Journal of Fluid Mechanics*, 865:281–302, 2019.
- [11] Giuseppe Da Prato and Arnaud Debussche. Control of the stochastic burgers model of turbulence. *SIAM Journal on Control and Optimization*, 37(4):1123–1149, 1999.
- [12] Scott J Moura and Hosam K Fathy. Optimal boundary control of reaction-diffusion partial differential equations via weak variations. *Journal of Dynamic Systems, Measurement, and Control*, 135(3):034501, 2013.
- [13] Donghui Yang and Jie Zhong. Optimal actuator location of the minimum norm controls for stochastic heat equations. *arXiv preprint arXiv:1710.06079*, 2017.
- [14] KB Lim. Method for optimal actuator and sensor placement for large flexible structures. *Journal of Guidance, Control, and Dynamics*, 15(1):49–57, 1992.
- [15] Tamara Nestorović and Miroslav Trajkov. Optimal actuator and sensor placement based on balanced reduced models. *Mechanical Systems and Signal Processing*, 36(2):271–289, 2013.
- [16] Dhanaraja Kasinathan and Kirsten Morris.  $h_\infty$ -optimal actuator location. *IEEE Transactions on Automatic Control*, 58(10):2522–2535, 2013.
- [17] Kevin K Chen and Clarence W Rowley.  $H_2$  optimal actuator and sensor placement in the linearised complex ginzburg-landau system. *Journal of Fluid Mechanics*, 681:241–260, 2011.
- [18] Krithika Manohar, J Nathan Kutz, and Steven L Brunton. Optimal sensor and actuator placement using balanced model reduction. *arXiv preprint arXiv:1812.01574*, 2018.
- [19] RO Grigoriev, MC Cross, and HG Schuster. Pinning control of spatiotemporal chaos. *Physical Review Letters*, 79(15):2795, 1997.
- [20] Subhrajit Sinha, Umesh Vaidya, and Rajeev Rajaram. Optimal placement of actuators and sensors for control of nonequilibrium dynamics. In *2013 European Control Conference (ECC)*, pages 1083–1088. IEEE, 2013.
- [21] Umesh Vaidya, Rajeev Rajaram, and Sambarta Dasgupta. Actuator and sensor placement in linear advection pde with building system application. *Journal of Mathematical Analysis and Applications*, 394(1):213–224, 2012.
- [22] Samuel Amstutz and Heiko André. A new algorithm for topology optimization using a level-set method. *Journal of computational physics*, 216(2):573–588, 2006.
- [23] Yiming Lou and Panagiotis D Christofides. Optimal actuator/sensor placement for nonlinear control of the kuramoto-sivashinsky equation. *IEEE Transactions on Control Systems Technology*, 11(5):737–745, 2003.
- [24] M Sajjad Edalatzadeh, Dante Kalise, Kirsten A Morris, and Kevin Sturm. Optimal actuator design for vibration control based on lqr performance and shape calculus. *arXiv preprint arXiv:1903.07572*, 2019.
- [25] G. Da Prato and J. Zabczyk. *Stochastic Equations in Infinite Dimensions*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 2008.
- [26] G. Da Prato and J. Zabczyk. *Stochastic Equations in Infinite Dimensions*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 2014.
- [27] Evangelos Theodorou. Nonlinear stochastic control and information theoretic dualities: Connections, interdependencies and thermodynamic interpretations. *Entropy*, 17(5):3352–3375, 2015.
- [28] Evangelos A Theodorou, George I Boutselis, and Kaivalya Bakshi. Linearly solvable stochastic optimal control for infinite-dimensional systems. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 4110–4116. IEEE, 2018.
- [29] George I Boutselis, Marcus A Pereira, Ethan N Evans, and Evangelos A Theodorou. Variational optimization for distributed and boundary control of stochastic fields.

- arXiv preprint arXiv:1904.02274*, 2019.
- [30] Grady Williams, Paul Drews, Brian Goldfain, James M. Rehg, and Evangelos A. Theodorou. Aggressive driving with model predictive path integral control. *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1433–1440, 2016.
- [31] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(Jul):2121–2159, 2011.
- [32] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [33] Gabriel J. Lord, Catherine E. Powell, and Tony Shardlow. *An Introduction to Computational Stochastic PDEs*. Cambridge Texts in Applied Mathematics. Cambridge University Press, 2014.
- [34] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.