

Reinforcement Learning for Infinite-Dimensional Systems

Wei Zhang

WEI.ZHANG@WUSTL.EDU

*Department of Electrical & Systems Engineering
Washington University in St. Louis
St. Louis, MO 63130, USA*

Jr-Shin Li

JSLI@WUSTL.EDU

*Department of Electrical & Systems Engineering
Division of Computational & Data Sciences
Division of Biology & Biomedical Sciences
Washington University in St. Louis
St. Louis, MO 63130, USA*

Editor: Nan Jiang

Abstract

Interest in reinforcement learning (RL) for large-scale systems, comprising extensive populations of intelligent agents interacting with heterogeneous environments, has surged significantly across diverse scientific domains in recent years. However, the large-scale nature of these systems often leads to high computational costs or reduced performance for most state-of-the-art RL techniques. To address these challenges, we propose a novel RL architecture and derive effective algorithms to learn optimal policies for arbitrarily large systems of agents. In our formulation, we model such systems as parameterized control systems defined on an infinite-dimensional function space. We then develop a moment kernel transform that maps the parameterized system and the value function into a reproducing kernel Hilbert space. This transformation generates a sequence of finite-dimensional moment representations for the RL problem, organized into a filtrated structure. Leveraging this RL filtration, we develop a hierarchical algorithm for learning optimal policies for the infinite-dimensional parameterized system. To enhance the algorithm's efficiency, we incorporate early stopping at each hierarchy, demonstrating the fast convergence property of the algorithm through the construction of a convergent spectral sequence. The performance and efficiency of the proposed algorithm are validated using practical examples in engineering and quantum systems.

Keywords: Reinforcement learning, ensemble systems, moment kernelization, spectral sequence convergence, control theory

1. Introduction

Reinforcement learning (RL), a prominent machine learning paradigm, has gained recognition as a powerful tool for intelligent agents to learn optimal policies. These policies guide the agents' decision-making processes toward achieving desired goals by maximizing the expected cumulative rewards. RL exhibits a wide range of applications across various domains, encompassing robotics, game-playing, recommendation systems, autonomous vehicles, and control systems. In recent years, learning optimal policies to manipulate the behavior of

large-scale systems of intelligent agents interacting with different environments has attracted increasing attention and is gradually becoming a recurrent research theme in the RL community. Existing and potential applications of “large-scale RL” include targeted coordination of robot swarms for motion planning in robotics (Becker and Bretl, 2012; Shahrokhi et al., 2018), desynchronization of neuronal ensembles with abnormal oscillatory activity for the treatment of movement disorders such as Parkinson’s disease, epilepsy, and essential tremor in neuroscience and brain medicine (Marks, 2005; Wilson, 2005; Shoeb, 2009; Zlotnik and Li, 2012; Li et al., 2013; Vu et al., 2024), and the robust excitation of spin ensembles for nuclear magnetic resonance (NMR) spectroscopy, imaging (MRI), and quantum computation in quantum science and technology (Glaser et al., 1998; Li et al., 2011; Dong et al., 2008; Chen et al., 2014).

The fundamental challenge in these RL tasks undoubtedly lies in the massive scale of agent systems and the dynamic environments in which the agents take actions. For instance, a neuronal ensemble in the human brain may comprise up to 10^{11} neurons (Herculano-Houzel, 2012; Ching and Ritt, 2013), and a quantum ensemble in NMR experiments consists of spins on the order of Avogadro’s number 10^{23} (Li, 2006; Cavanagh et al., 2010; Li, 2011). Although, mathematically speaking, these systems are composed of finitely many agents, it is more appropriate to treat them as infinite populations. This is because their massive scale makes it infeasible to specifically identify and make decisions for each individual agent in these populations. These restrictions particularly imply that such policy learning tasks exceed the capabilities of classical multi-agent RL methods, which focus on learning joint policies that provide each agent with a customized action based on the states of all the other agents (Littman, 1994; Foerster et al., 2016; Gupta et al., 2017; Zhang et al., 2021b; Albrecht et al., 2024). Consequently, all the agents are required to take the same action, necessitating the learning and implementation of policies at the population level. In addition, natural models describing the dynamic behavior of these intelligent agents typically take the form of continuous-time deterministic control systems. Learning optimal policies for this category of systems also lies outside the scope of most state-of-the-art and benchmark RL algorithms, notably Trust Region Policy Optimization (TRPO) and Proximal Policy Optimization (PPO) (Schulman et al., 2017, 2015), which are based on Markov decision processes (MDPs) and operate as discrete-time stochastic control processes. This, in turn, underscores the urgent demand for a more inclusive RL framework that accommodates policy learning tasks for infinite populations of agents in a continuous-time deterministic setting.

Our contributions. This work is devoted to developing a novel RL architecture that enables the derivation of effective algorithms to learn optimal policies for population systems consisting of infinitely many intelligent agents interacting with heterogeneous dynamic environments. We formulate such an agent system as a parameterized control system, where each individual system indexed by a specific parameter value represents the environment of an agent in the population. We then evolve this parameterized system on an infinite-dimensional function space and carry out a functional setting for RL of this infinite-dimensional dynamical system. The primary tool we develop to tackle this RL problem is the moment kernel transform. This transformation maps the parameterized system and the value function to a control system and a value function defined on a reproducing kernel Hilbert space (RKHS) consisting of moment sequences, thus yielding a kernel parameterization of the

RL problem. The use of moment sequences for this kernelization directly enables finite-dimensional truncation representations of the infinite-dimensional RL problem. We then organize these representations into a filtrated structure with respect to the truncation order. Leveraging this RL filtration, we develop a hierarchical policy learning algorithm, wherein each hierarchy consists of an RL problem for a finite-dimensional truncated moment-kernelized system. To enhance the computational efficiency of the proposed algorithm, we also develop early stopping criteria for the finite-dimensional RL problem in each hierarchy and prove convergence of the hierarchical algorithm with early-stopped hierarchies in terms of spectral sequences. The performance and efficiency of the proposed hierarchical policy learning algorithm are demonstrated using examples arising from practical applications. The contributions of our work are summarized as follows.

- Formulation of an infinite population of intelligent agents as a parameterized control system defined on an infinite-dimensional function space.
- Development of the moment kernel transform that gives rise to a kernel parameterization of RL problems for parameterized systems in terms of moment sequences in an RKHS.
- Design of a filtrated RL algorithm for learning optimal policies of parameterized systems with convergence guarantees.
- Exploration of early stopping criteria for each hierarchy in the proposed hierarchical algorithm and proof of the spectral sequence convergence of the hierarchical algorithm with early-stopped hierarchies.

Related works. Learning control policies to coordinate large populations of intelligent agents using RL approaches has attracted increasing attention in recent years. Significant achievements have been recognized in the field of control theory and engineering. Prominent examples include control of multi-agent systems (Buşoniu et al., 2010; Heredia and Mou, 2019; Jiang et al., 2021; Zhang et al., 2021b; Albrecht et al., 2024) and many-body quantum systems (Dong et al., 2008; Lamata, 2017; Bukov et al., 2018; Haug et al., 2021), as well as planning and regulation for multi-robot systems (Matarić, 1997; Long et al., 2018; Yang et al., 2020a; Yang and Gu, 2004; Lu et al., 2024). Furthermore, large-scale RL has also gained great success in various emerging applications that are generally considered to be outside the traditional control-theoretic scope, notably, strategic gaming (Silver et al., 2017; Sil, 2018; Vinyals et al., 2019; OpenAI et al., 2019; Schrittwieser et al., 2020; Kaiser et al., 2020) and human-level decision making (Mnih et al., 2015; Liu et al., 2022; Baker et al., 2020).

One of the most active research focuses regarding large-scale and high-dimensional policy learning problems is deep RL, which incorporates deep learning techniques, particularly the use of deep neural networks, into RL algorithms (Bertsekas and Tsitsiklis, 1996; François-Lavet et al., 2018; Bellemare et al., 2020; Nakamura-Zimmerer et al., 2021; Le et al., 2022; Sarang and Poullis, 2023). Despite the advantage of exceptional generalizability, training deep neural networks to tackle large-scale and/or high-dimensional learning problems is widely known to suffer from the notorious curse of dimensionality, leading to expensive computational costs and scalability issues. Various approaches have been proposed to mitigate the impact of these phenomena, primarily including the development of distributed and multi-agent RL

algorithms (Heredia and Mou, 2019; Heredia et al., 2020; Yazdanbakhsh et al., 2020; Heredia et al., 2022; Xie et al., 2024) and the search for compact representations of high-dimensional measurement data. Notable among these are the successor representation (Momennejad et al., 2017), latent space representation (Gelada et al., 2019), contrastive unsupervised representation (Laskin et al., 2020), and invariant representation (Zhang et al., 2021a). These works have achieved considerable success in learning optimal policies for agent populations on a scale ranging from tens to thousands.

To further push the boundaries of RL towards addressing larger agent populations, tools in mean-field theory - particularly mean-field games and control - have drawn increasing attention and have been adopted in the RL setting in recent years. Instead of learning a policy for each individual agent in a population, mean-field RL focuses on the mean-field approximation of the population, where each collection of interacting agents is replaced by a single agent representing their averaged behavior (Yang et al., 2018; Subramanian and Mahajan, 2019; Laurière et al., 2022; Pásztor et al., 2023; Bensoussan et al., 2013; Carmona et al., 2019; Fu et al., 2020; Carmona et al., 2020). However, the prerequisite of the mean-field approximation, arising from the fundamental principles of statistical physics, places the primary focus of mean-field RL on populations of identical agents (Pathria and Beale, 2021). In this context, policy learning for arbitrarily large, and in the limit, infinite populations of heterogeneous intelligent agents, as considered in this work, remains under-explored with sparse literature in the RL community. On the other hand, the formulation of RL problems over infinite-dimensional spaces has only been proposed in the stochastic setting, aiming to learn feedback control policies for stochastic partial differential equation systems using variational optimization methods (Evans et al., 2020).

The major technical tool developed in this work to overcome the curse of dimensionality is the moment kernel transform, which is inspired by the method of moments. This method was developed by the Russian mathematician P. L. Chebyshev in 1887 to prove the law of large numbers and the central limit theorem (Mackey, 1980). Since then, the method has been extensively studied under different settings, notably the Hausdorff, Hamburger, and Stieltjes moment problems (Hamburger, 1920, 1921a,b; Hausdorff, 1923; Stieltjes, 1993). The most general formulation in modern terms was proposed by the Japanese mathematician Kōsaku Yosida (Yosida, 1980). Recently, the method of moments was introduced to control theory for establishing dual representations of ensemble systems (Narayanan et al., 2024; Li et al., 2025) and in machine learning-aided medical decision making as a feature engineering technique (Yu et al., 2023). These prior works lay the foundation for the development of the moment kernel transform in this work.

In addition to establishing the filtrated RL architecture, the moment kernel transform also gives rise to a reduced kernel representation of (infinite) agent populations over an RKHS. It is widely known that RKHS theory forms the building blocks for kernel methods in machine learning, notably support vector machines (SVMs) and kernel principal component analysis (kernel PCA) (Hastie et al., 2009; Paulsen and Raghupathi, 2016). In the context of RL, elements of RKHSs are commonly used as function approximators for MDPs, through which the learning targets, including policies, value functions, and/or transition maps, are estimated in terms of linear combinations of reproducing kernels (Lever and Stafford, 2015; Yang and Wang, 2019; Yang et al., 2020b; Koppel et al., 2021). This kernel approximation is typically enabled by imposing the condition that these learning targets belong to an RKHS

consisting of functions defined on the state space of the MDPs. In our work, we explore the use of RKHS-theoretic techniques within a functional setup rather than the traditional MDP setup, thus relaxing the condition on the learning targets. More importantly, along with enabling kernel approximation, the developed moment kernel transform also serves as a model reduction machine for arbitrarily large agent populations.

2. Policy Learning for Infinite Dynamic Populations

This section is primarily dedicated to establishing a general formulation for RL of arbitrarily large populations of intelligent agents, which interact with heterogeneous environments in continuous time and are regulated by policies taking values in continuous spaces. We begin by demonstrating the challenges associated with such RL tasks, and then introduce our formulation that models these populations as continuous-time parametrized control systems defined on infinite-dimensional function spaces. We subsequently delve into imposing conditions that guarantee the existence of optimal policies for RL problems involving these infinite-dimensional control systems over function spaces.

2.1 Challenges to reinforcement learning of large-scale population agents

Learning control policies for a large-scale population of intelligent agents, effectively a continuum in the limit, described by continuous-time deterministic dynamical systems, presents a significant challenge to RL. The primary obstacle is the enormous population size, which forces RL algorithms to operate in high-dimensional spaces. This inevitably leads to the curse of dimensionality, resulting in high computational costs and reduced learning accuracy (Bellman et al., 1957; Bellman, 1961; Sutton and Barto, 2018). Additionally, the continuous-time deterministic system formulation is inconsistent with the setting of most state-of-the-art RL algorithms, which are typically based on MDPs modeled by discrete-time stochastic control processes.

To illustrate these challenges with a concrete example, we consider a population system consisting of N dynamic agents regulated by a common control policy, given by

$$\frac{d}{dt}x_i(t) = a_i x_i(t) + u(t), \quad i = 1, \dots, n, \quad (1)$$

where $x_i(t) \in \mathbb{R}$ is the state of agent i , $u(t) \in \mathbb{R}$ is the control policy, and $a_i \in \mathbb{R}$ for all i . The entire size- n population can be represented as an n -dimensional system, $\frac{d}{dt}x(t) = A_n x(t) + B_n u(t)$, where $x(t) = [x_1(t), \dots, x_n(t)]' \in \mathbb{R}^n$ is the population state with $'$ denoting the transpose of vectors (and matrices), $A_n \in \mathbb{R}^{n \times n}$ is the diagonal matrix with the (i, i) -entry given by a_i for all $i = 1, \dots, n$, and $B_n \in \mathbb{R}^n$ is a vector of ones. To put the analysis in a simplified setting, we choose $-1 = a_0 < a_1 < \dots < a_n = 1$ to be the uniform partition of the interval $[-1, 1]$, i.e., $a_i = -1 + 2(i - 1)/(n - 1)$ for each $i = 1, \dots, n$. We aim to learn the infinite-time horizon linear quadratic regulator (LQR) with the state-value function (cumulative reward or cost-to-go), $V(x(t)) = \int_t^\infty e^{-2.5t} \left[\frac{1}{n} \sum_{i=1}^n x_i^2(t) + u^2(t) \right] dt = \int_t^\infty e^{-2.5t} \left[\frac{1}{n} x'(t)x(t) + u^2(t) \right] dt$. It is well-known that the LQR value function $V^*(x(t)) = \inf_u V(x(t))$ is in the quadratic form $V^*(x(t)) = x'(t)Qx(t)$, parameterized by a positive definite matrix $Q \in \mathbb{R}^{n \times n}$ (Brockett, 2015). Therefore, the value iteration stands out as the

prime algorithm to tune the $n(n+1)/2$ training parameters in Q to learn the LQR policy and value function (Sutton and Barto, 2018; Bertsekas, 2019).

Curse of dimensionality. In the simulation, we varied the system dimension n from 2 to 20. For each n , we independently ran the value iteration 5 times with random initial conditions. The simulation results are shown in Figure 1. In particular, Figure 1a shows the number of training parameters (top panel) and the average computational time over the 5 runs of the value iteration versus N (bottom panel). From these plots, we observe dramatic increases in both quantities. In addition to these common phenomena associated with the curse of dimensionality, we also observed convergence issues in this high-dimensional system policy learning problem.

Convergence issue. Note that by the definition of Riemann integrals (Rudin, 1976), for each t , the term $\frac{1}{n} \sum_{i=1}^n x_i^2(t)$ in the state-value function V is essentially a Riemann sum of the real-value function $\frac{1}{2}x(t, \cdot)$ over the interval $[-1, 1]$. Therefore, by the Dominated Convergence theorem (Folland, 2013), V possesses the convergence property,

$$V(x(t)) = \int_t^\infty e^{-2.5t} \left[\frac{1}{n} \sum_{i=1}^n x_i^2(t) + u^2(t) \right] dt \rightarrow \int_t^\infty \frac{e^{-2.5t}}{2} \left[\int_{-1}^1 x^2(t, \beta) d\beta + u^2(t) \right] dt \quad (2)$$

as $N \rightarrow \infty$, where $x(t, \beta)$ satisfies the linear parameterized system on \mathbb{R} , given by

$$\frac{d}{dt}x(t, \beta) = \beta x(t, \beta) + u(t), \quad \beta \in \Omega = [-1, 1]. \quad (3)$$

To see that the value function also inherits this convergence property, we notice that the state-value function is continuous, and hence upper semi-continuous, for any control policy u . This implies that the values function $V(x(t)) = \inf_u V(x(t))$ is upper semi-continuous (Folland, 2013). Together with the non-negativity of V , we obtain the convergence of the state-value function, $\lim_{N \rightarrow \infty} \inf_u V(x(t)) = \inf_u \lim_{N \rightarrow \infty} V(x(t)) = \inf_u \int_t^\infty \frac{e^{-2.5t}}{2} \left[\int_{-1}^1 x^2(t, \beta) d\beta + u^2(t) \right] dt$ (see Theorem 3 below for the justification that the right-hand side is well-defined) (Folland, 2013). Let V_n^* and u_n^* denote the value function and LQR policy learned from the n -dimensional system. As shown in Figure 1b, neither $\|V_n^* - V_{n-1}^*\| = \sup_{t \geq 0} |V_n^*(x^*(t)) - V_{n-1}^*(x^*(t))|$ nor $\|u_n^* - u_{n-1}^*\| = \sup_{t \geq 0} |u_n^*(t) - u_{n-1}^*(t)|$ shows a trend of converging to 0, where $x^*(t)$ denotes the optimal trajectory. This particularly fails to verify the convergence of the value function as illustrated in (2).

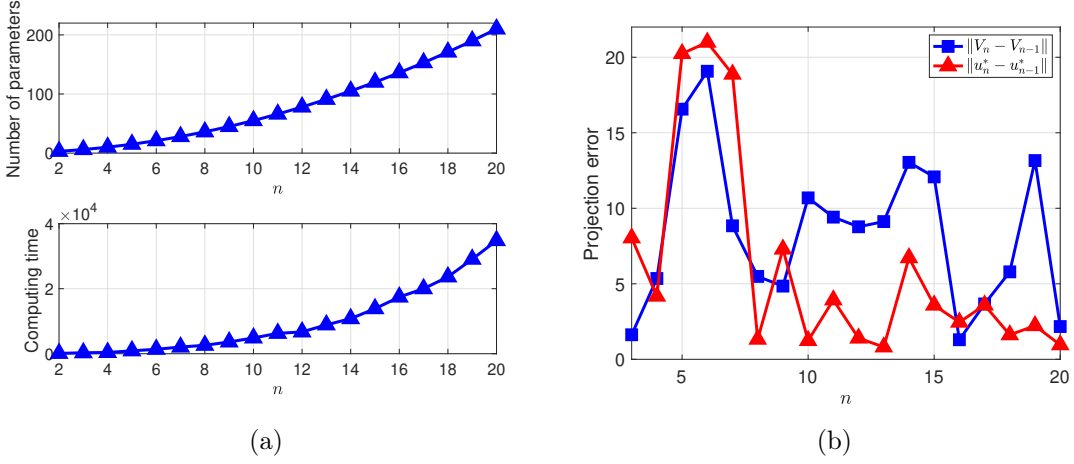


Figure 1: Illustration of the challenges to RL for high-dimensional systems. The standard value iteration is applied to learn the LQR policy and value function for an n -dimensional time-invariant deterministic linear system, with N ranging from 2 to 20. (a) shows the number of training parameters (top) and average computing time (bottom) with respect to the system dimension n . (b) plots the average distance between successive terms in the sequences of the learned value functions (blue) and LQR policies (red) with respect to n .

To further elaborate on this convergence issue, we treat the parameterized system in (3) as a continuum ensemble of linear systems indexed by $\beta \in [-1, 1]$. The finite ensemble in (1) can be seen as essentially a size- n sample of this infinite ensemble. From this perspective, the convergence issue arises due to the ill-posed problem stemming from the sampling (discretizing) of the continuum ensemble. Specifically, the optimal policy for a size- n sample may be highly suboptimal for a size- m ensemble for $m > n$, so that u_m is significantly different from u_n . To carry out a quantitative analysis in the simplest setting, we consider the case of learning a policy that steers the entire infinite ensemble to the origin 0 with minimal energy. Specifically, the cumulative reward function is given by $V(t, x(t)) = \int_t^1 u^2(t) dt$. The minimum energy policy for the size- n sample is given by $u_n(t) = B_n' e^{-A_n t} W_n^{-1}(0, 1) x_{0n}$ so that $V_n(t, x(t)) = x_{0n}' W_n^{-1}(t, 1) x_{0n}$, where $W_n(t_1, t_2) = \int_{t_1}^{t_2} e^{-A_n s} B_n B_n' e^{-A_n' s} ds$ for $0 \leq t_1 \leq t_2 \leq 1$ is referred to as the *controllability Gramian*, and x_{0n} is the initial condition of the size- n ensemble system (Brockett, 2015; Liberzon, 2012). It is not hard to see that the analytic properties of V_n and u_n primarily depend on those of W_n . Specifically, the (i, j) -entry of $W_n(t_1, t_2)$ is given by $\frac{e^{-(a_i + a_j)t_1} - e^{-(a_i + a_j)t_2}}{a_i + a_j}$. Its trace, $\text{tr}(W_n(t_1, t_2)) = \sum_{i=1}^n \frac{e^{-2a_i t_1} - e^{-2a_i t_2}}{2a_i}$, satisfies $\frac{1 - e^{-2}}{2} \leq \frac{1}{n} \text{tr}(W_n(t_1, t_2)) \leq \frac{e^2 - 1}{2}$. Hence, all the eigenvalues of $W(t_1, t_2)$ are bounded below and above by $\lambda_{\min} = \frac{1 - e^{-2}}{2}$ and $\lambda_{\max} = \frac{e^2 - 1}{2}$, respectively. As a result, for any different sample sizes m and n , there are initial conditions x_{0n} and x_{0m} such that $\sup_t |V_n(t, x(t)) - V_m(t, x(t))| \geq \lambda_{\min}^{-1} - \lambda_{\max}^{-1} = 2$ and $\sup_t |u_n(t) - u_m(t)| \geq 2$, highlighting the convergence issue.

To overcome the presented challenges, we propose a new kernel parameterization technique that will serve as a foundational element for developing a novel RL architecture. To facilitate this, we will first introduce a principled RL formulation that accounts for large-scale population systems, regardless of their size.

2.2 Reinforcement learning for parameterized systems on function spaces

The parameterized representation in (3) of the “linear agents” population in the limiting case inspires the modeling of agent populations of any size as parameterized differential equation systems, also referred to as *ensemble systems*, of the form

$$\frac{d}{dt}x(t, \beta) = F(t, \beta, x(t, \beta), u(t)), \quad (4)$$

where β is the system parameter taking values in $\Omega \subseteq \mathbb{R}^d$, $x(t, \beta) \in M$ is the state of the system indexed by the particular value of β (the environmental state of agent β) in the population with $M \subseteq \mathbb{R}^n$ being a differentiable manifold, $u(t) \in \mathbb{R}^r$ is the control policy, and $F(t, \beta, \cdot, u(t))$ is a (time-varying) vector field on M for each $\beta \in \Omega$ and $u(t) \in \mathbb{R}^r$, characterizing the environment of the agent β . Similarly, as motivated by the Riemann sum convergence illustrated in (2), we associate the parametrized ensemble system in (4) with a state-value function in the integral form, given by

$$V(t, x_t) = \int_{\Omega} \left[\int_t^T r(s, x(s, \beta), u(s)) ds + K(T, x(T, \beta)) \right] d\beta, \quad (5)$$

where $x_t(\cdot) \doteq x(t, \cdot)$, and $r(s, x(s, \beta), u(s))$ and $K(T, x(T, \beta))$ represent the running and terminal costs of agent β , respectively. One of the major advantages of the parameterized system formulation is the flexibility to model agent populations of arbitrarily large size, extending to a continuum of agents when the parameter space Ω has an uncountable cardinality, as in the case of (3). However, due to practical limitations on sensing capability and computing power, it is impossible to collect comprehensive measurement data documenting the state and reward information for all the agents in such a massive-scale agent population. This constraint particularly hinders the design and implementation of feedback control policies, placing the relevant policy learning problem beyond the scope of many existing RL and optimal control methods.

To develop a new RL paradigm suitable for this type of policy learning tasks, we reframe the parameterized system in (4) and its state-value function in (5) from a different perspective. Specifically, the ensemble state of the parameterized system can be seen as a function $x(t, \cdot) : \Omega \rightarrow M$, such that the system is evolving on a space $\mathcal{F}(\Omega, M)$ of M -valued functions defined on Ω . Thus, the state-value function is a functional on $\mathcal{F}(\Omega, M)$, essentially characterizing the “average” of the cumulative rewards of all the agents over the entire population. This functional viewpoint, in turn, situates the policy learning tasks within the function space $\mathcal{F}(\Omega, M)$. When Ω is an infinite set, $\mathcal{F}(\Omega, M)$ becomes an infinite-dimensional manifold. Therefore, to ensure the solvability of these policy learning tasks, some regularity conditions on the system dynamics F , immediate reward r , and terminal cost K , more stringent than the canonical setup in RL, are required.

Assumption S1. (Boundedness of control policies) The ensemble control policy $u : [0, T] \rightarrow \mathbb{R}^r$ is a measurable function and takes values in a compact subset U of \mathbb{R}^r .

Assumption S2. (Lipschitz continuity of system dynamics) The vector field $F : \mathbb{R} \times \Omega \times M \times U \rightarrow TM$ is continuous in all variables and Lipschitz continuous in $z \in M$ uniformly for

$(t, \beta, a) \in [0, T] \times \Omega \times U$. Specifically, there exists a constant C (independent of $t \in [0, T]$, $\beta \in \Omega$, and $a \in U$) such that $|(\varphi_* F|_V)(t, \beta, \varphi(p), a) - (\varphi_* F|_V)(t, \beta, \varphi(q), a)| \leq C|\varphi(p) - \varphi(q)|$ for any $p, q \in V$ and coordinate chart (V, φ) on M . Here, TM denotes the tangent bundle of M , $\varphi_* F|_V$ denotes the pushforward of $F|_V$, the restriction of the vector field F on $V \subseteq M$, and $|\cdot|$ denotes a norm on \mathbb{R}^n .

According to the theory of ordinary differential equations, Assumption S2 guarantees that, driven by any admissible control policies, each individual system in the ensemble indexed by β in (4) has a unique and Lipschitz continuous solution $t \mapsto x(t, \beta)$ (Arnold, 1978; Lang, 1999). Correspondingly, at the population level, where β varies over Ω , the ensemble system has a unique solution in $\mathcal{F}(\Omega, M)$, given by $t \mapsto x(t, \beta)$.

Assumption C1. (Integrability of the state-value function) There exists an ensemble control policy $u \in \mathcal{U}$ such that $\int_{\Omega} \int_0^T |r(t, x(t, \beta), u(t))| dt d\beta + \int_{\Omega} |K(T, x(T, \beta))| d\beta < \infty$, where $x(t, \cdot) \in \mathcal{F}(\Omega, M)$ is the solution of the ensemble system in (4) driven by u .

Assumption C2. (Lipschitz continuity of private costs) Both the running cost $r : \mathbb{R} \times M \times U \rightarrow \mathbb{R}$ and terminal cost $K : \mathbb{R} \times M \rightarrow \mathbb{R}$ are continuous functions in all variables and Lipschitz continuous in $z \in M$ for any $(t, a) \in [0, T] \times U$.

In the sequel, these regularity assumptions will be exploited to prove the existence of a solution to the functional policy learning problem formulated in (4) and (5). In other words, the value function $V^*(t, x_t) = \inf_u V(t, x_t)$ is a well-defined real-valued function on $[0, T] \times \mathcal{F}(\Omega, M)$. To emphasize the dependence of the state-value function V on the policy u , we denote the total cost by $V(0, x_0) = J(u)$, which defines a function $J : \mathcal{U} \rightarrow \mathbb{R}$, referred to as the cost functional, with $\mathcal{U} = \{u : [0, T] \rightarrow \mathbb{R}^r \mid u(t) \in U\}$ being the space of admissible policies. The solvability of the learning problem then hinges on proving the compactness of \mathcal{U} and continuity of J .

Lemma 1 (Compactness of admissible policies) *The space of admissible policies \mathcal{U} is compact.*

Proof Topologically, \mathcal{U} , equipped with the topology of pointwise convergence, is the product space $\prod_{t \in [0, T]} U$ under the product topology. Because $U \subset \mathbb{R}^m$ is compact by Assumption S1, the compactness of \mathcal{U} directly follows from Tychonoff's theorem (Munkres, 2000). ■

Lemma 2 (Sequential continuity of the cost functional) *The cost functional $J : \mathcal{U} \rightarrow \mathbb{R}$ is sequentially continuous, that is, $J(u_k) \rightarrow J(u)$ for any admissible policy sequence $(u_k)_{k \in \mathbb{N}}$ such that $u_k \rightarrow u$ in \mathcal{U} .*

Proof See Appendix A. ■

Theorem 3 (Existence of optimal policies) *Given a parameterized ensemble system defined on the function space $\mathcal{F}(\Omega, M)$ as in (4) satisfying Assumptions S1 and S2, and the*

state-value function defined in (5) satisfying Assumptions C1 and C2, the value function $V^* : [0, T] \times \mathcal{F}(\Omega, M) \rightarrow \mathbb{R}$, given by $V^*(t, x_t) = \inf_u V(t, x_t)$, is well-defined. This implies that an optimal policy exists.

Proof By the uniqueness of the solution to the parameterized ordinary differential equation in (4), guaranteed by Assumptions S1 and S2, if there is a policy $u^* \in \mathcal{U}$ that minimizes J , then it necessarily minimizes $V(t, x_t)$ for all $t \in [0, T]$ as well. In addition, we have $J(u^*) = \min_{u \in \mathcal{U}} J(u) = V^*(0, x_0)$ by the definition of J , which is unique as the infimum (greatest lower bound) of the set $J(\mathcal{U}) \subset \mathbb{R}$ (Rudin, 1976). As a result, $V^*(0, x_0)$, and hence $V^*(t, x_t)$ for each $t \in [0, T]$, has a definite real value, indicating that V^* is a well-defined (i.e., single-valued) real-valued function on $[0, T] \times \mathcal{F}(\Omega, M)$. Therefore, it suffices to show the existence of $u^* \in \mathcal{U}$.

Without loss of generality, we assume that the integrability condition in Assumption C1 is satisfied for all $u \in \mathcal{U}$. We then show that the range $J(\mathcal{U})$ of the cost functional J is a compact subspace of \mathbb{R} , equivalently, any sequence in $J(\mathcal{U})$ has a convergent subsequence, because of $J(\mathcal{U}) \subseteq \mathbb{R}$ (Munkres, 2000). To show this, we pick an arbitrary sequence $(J_k)_{k \in \mathbb{N}}$ in $J(\mathcal{U})$ and claim that any policy sequence $(u_k)_{k \in \mathbb{N}}$ satisfying $J_k = J(u_k)$ for all $k \in \mathbb{N}$ has an accumulation point $u \in \mathcal{U}$. To see this, we define $T_k = \{u_n : n \geq k\}$. Then, $\bigcap_{k \in \mathbb{N}} \bar{T}_k \neq \emptyset$ holds, where \bar{T}_n denotes the closure of T_n . Otherwise, $\{\mathcal{U} \setminus \bar{T}_k\}_{k \in \mathbb{N}}$ forms an open cover of \mathcal{U} , which has a finite subcover, say $\mathcal{U} = \bigcup_{k=1}^N (\mathcal{U} \setminus \bar{T}_k)$, due to the compactness of \mathcal{U} shown in Lemma 1, leading to the contradiction $\emptyset = \mathcal{U} \setminus (\bigcap_{k=1}^N \bar{T}_k) = \bigcup_{k=1}^N (\mathcal{U} \setminus \bar{T}_k) = \bigcup_{k=1}^N \bar{T}_k = \bar{T}_N$. Now, let $u \in \bigcap_{k \in \mathbb{N}} \bar{T}_k$ and W be a neighborhood of u . Then $W \cap T_k \neq \emptyset$ for every k , i.e., W contains some u_k for any arbitrarily large k , and hence u is necessarily an accumulation point of the sequence $(u_k)_{k \in \mathbb{N}}$. The sequential continuity of J , proved in Lemma 2, implies that J necessarily maps accumulation points of policy sequences to accumulation points in $J(\mathcal{U})$. As a result, there is a subsequence of $(J_k)_{k \in \mathbb{N}}$ that converges to $J(u)$, showing the sequential compactness, and hence compactness, of $J(\mathcal{U})$.

The compactness of $J(\mathcal{U})$ particularly implies that $J^* = \inf J(\mathcal{U}) = \inf_{u \in \mathcal{U}} J(u) \in J(\mathcal{U})$. Therefore, there exists $u^* \in \mathcal{U}$ such that $J(u^*) = J^*$, meaning u^* is an optimal policy. ■

Remark 4 (Topological actor-critic algorithm) Although Theorem 3 verifies the solvability of the policy learning problem over the infinite-dimensional function space from a theoretical perspective, the main idea of the employed topological argument aligns closely with the actor-critic algorithm in RL. Specifically, the “critic” J constantly evaluates the “actor” u to iteratively improve its performance, generating a sequence of cost J_k converging to the minimal cost J^* . The corresponding policy u^* with $J^* = J(u^*)$ is then a desired optimal policy. Moreover, importantly, approaching u^* through a cost sequence instead of a policy sequence deliberately avoids a technical issue. Although \mathcal{U} is compact, a policy sequence u_k in \mathcal{U} may not contain any convergent subsequence since \mathcal{U} is not a first-countable space (Munkres, 2000).

3. Reinforcement Learning for Parameterized Systems via Moment Parameterization

In this section, we will focus on developing an RL framework for learning optimal policies for parameterized ensemble systems defined on an infinite-dimensional function space. Our initial step, which is also essential to most learning problems, is to explore an appropriate parameterization for the learning targets. To this end, we will introduce a moment kernel transform, which generates kernel representations of parameterized systems and state-value functions over a reproducing kernel Hilbert space (RKHS).

3.1 Moment kernelization of parameterized systems and value functions

Our theoretical development is based on leveraging and extending the method of moments in functional analysis and probability theory (Mackey, 1980). The central idea is to represent the time-varying state functions of a parameterized system as time-dependent sequences of real numbers. To formalize this idea, we impose the following assumption:

Assumption K1. The state space $\mathcal{F}(\Omega, M)$ of the parameterized ensemble system in (4), given by $\frac{d}{dt}x(t, \beta) = F(t, \beta, x(t, \beta), u(t))$ with $\beta \in \Omega \subset \mathbb{R}^d$, is a Hilbert space \mathcal{H} contained in $L^2(\Omega, \mathbb{R}^n)$, the space of \mathbb{R}^n -valued square-integrable functions defined on Ω .

Ensemble moments and moment kernel transform. To motivate our idea of dynamic moment kernelization, we consider the scalar-valued parameterized system defined on the Hilbert space $\mathcal{H}_0 \subset L^2(\Omega, \mathbb{R})$; namely, the ensemble state $x(t, \cdot) \doteq x_t(\cdot) \in \mathcal{H}_0$. Because $L^2(\Omega, \mathbb{R})$ is separable, \mathcal{H}_0 is also separable as a linear subspace of \mathcal{H} (Yosida, 1980). Hence, \mathcal{H}_0 possesses a countable orthonormal basis, denoted $\{\Phi_k\}_{k \in \mathbb{N}}$. We define the k^{th} *moment* of the parameterized system for each $k \in \mathbb{N}$ with respect to Φ_k as

$$m_k(t) = \langle \Phi_k, x_t \rangle, \quad (6)$$

where $\langle \cdot, \cdot \rangle : \mathcal{H}_0 \times \mathcal{H}_0 \rightarrow \mathbb{R}$ denotes the inner product in \mathcal{H}_0 . Then, the moment sequence $m(t) \in \mathcal{M}_0$ associated with $x_t \in \mathcal{H}_0$ is denoted by $m(t) = (m_k(t))_{k \in \mathbb{N}}$, with \mathcal{M}_0 being the space of all moment sequences, referred to as the *moment space*. In addition, the inner product on \mathcal{H}_0 , as a subspace of $L^2(\Omega, \mathbb{R})$ illustrated in Assumption K1, is specifically given by $m_k(t) = \langle \Phi_k, x_t \rangle = \int_{\Omega} \Phi_k(\beta) x_t(\beta) d\beta$. The choice of $\{\Phi_k\}_{k \in \mathbb{N}}$ as an orthonormal basis yields $\sum_{k \in \mathbb{N}} |m_k(t)|^2 = \|m(t)\|_{\mathcal{M}_0}^2 = \|x_t\|_{\mathcal{H}_0}^2 = \int_{\Omega} |x_t(\beta)|^2 d\beta < \infty$ by Parseval's identity (Folland, 2013), where $\|\cdot\|_{\mathcal{M}_0}$ and $\|\cdot\|_{\mathcal{H}_0}$ denote the norms on \mathcal{M}_0 and \mathcal{H}_0 , respectively. This implies that \mathcal{M}_0 is contained in the ℓ^2 -space, consisting of square-summable sequences. Essentially, $m_k(t)$ is the k^{th} Fourier coefficient of the function $x_t \in \mathcal{H}_0$, and thus the moment sequence $m(t)$ provides a coordinate representation of x_t with respect to the basis $\{\Phi_k\}_{k \in \mathbb{N}}$.

In the general case where $n > 1$, the state space \mathcal{H} of the parameterized system admits a decomposition as a direct sum of n copies of \mathcal{H}_0 , i.e., $\mathcal{H} = \mathcal{H}_0 \oplus \cdots \oplus \mathcal{H}_0$. Equivalently, each component x_t^i of the ensemble state $x_t = (x_t^1, \dots, x_t^n)$ is an element of \mathcal{H}_0 . As a result, the definition of moments in (6) can be extended to $x_t \in \mathcal{H}$ in a component-wise manner as $m_k(t) = \langle \Phi_k, x_t \rangle = (\langle \Phi_k, x_t^1 \rangle, \dots, \langle \Phi_k, x_t^n \rangle)$. Then, $m_k(t)$ becomes an \mathbb{R}^n -valued sequence so that the moment space $\mathcal{M} = \mathcal{M}_0 \oplus \cdots \oplus \mathcal{M}_0$ is a Hilbert subspace of $\ell^2 \oplus \cdots \oplus \ell^2$, the n -fold direct sum of ℓ^2 -spaces. More importantly, this indicates that \mathcal{M} is an RKHS as shown below.

Proposition 5 *The moment space \mathcal{M} is an RKHS on \mathbb{N} , and the moment kernel transform $\mathcal{K} : \mathcal{H} \rightarrow \mathcal{M}$, given by $x_t \mapsto m(t)$, is an isometric isomorphism between the Hilbert spaces \mathcal{H} and \mathcal{M} .*

Proof Following the analysis above, we know that the moment space \mathcal{M}_0 of functions in \mathcal{H}_0 is contained in ℓ^2 . It remains to show that $\ell^2 \subseteq \mathcal{M}_0$. To this end, we pick any $m(t) \in \ell^2$. According to the Pythagorean theorem, we have $|\sum_{k=0}^{\infty} m_k(t) \Phi_k| = \sum_{k=0}^{\infty} |m_k(t)|^2 < \infty$, implying that the partial sum $\sum_{k=0}^N m_k(t) \Phi_k$ forms a Cauchy sequence. Therefore, $x_t = \sum_{k=0}^{\infty} m_k(t) \Phi_k$ is a well-defined element (function) in \mathcal{H}_0 and satisfies $\langle \Phi_k, x_t \rangle = m_k(t)$, yielding $\ell^2 \subseteq \mathcal{M}_0$. This concludes that $\mathcal{M}_0 = \ell^2$. The restriction of the moment transform to each component of $x_t \in \mathcal{H}$ is thus an isometric isomorphism from \mathcal{H}_0 to \mathcal{M}_0 . Together with the fact that $\mathcal{H} = \mathcal{H}_0 \oplus \cdots \oplus \mathcal{H}_0$, we obtain $\mathcal{M} = \mathcal{M}_0 \oplus \cdots \oplus \mathcal{M}_0 = \ell^2 \oplus \cdots \oplus \ell^2$, and the moment kernel transform from \mathcal{H} to \mathcal{M} is an isometric isomorphism as well.

Lastly, to show that \mathcal{M} is an $(\mathbb{R}^n\text{-valued})$ RKHS on \mathbb{N} , it suffices to demonstrate that the point evaluation map $E_k : \mathcal{M} \rightarrow \mathbb{R}^n$, given by $m(t) \mapsto m_k(t)$, is bounded (Paulsen and Raghupathi, 2016). This follows from the estimate,

$$|m_k(t)|^2 = \sum_{i=1}^n |\langle \Phi_k, x_t^i \rangle|^2 \leq \sum_{i=1}^n \|\Phi_k\|_{\mathcal{H}_0}^2 \|x_t^i\|_{\mathcal{H}_0}^2 = \sum_{i=1}^n \|x_t^i\|_{\mathcal{H}_0}^2 = \|x_t\|_{\mathcal{H}}^2 = \|m(t)\|_{\mathcal{M}}^2,$$

where $|\cdot|$, $\|\cdot\|_{\mathcal{H}}$, and $\|\cdot\|_{\mathcal{M}}$ denote the norms on \mathbb{R}^n , \mathcal{H} , and \mathcal{M} , respectively. ■

Moment kernelization of ensemble systems. Having kernelized the ensemble state of the parametrized system, the next step is to exploit the kernelized state and the moment transform to kernelize the system dynamics, i.e., the temporal evolution of the ensemble state over \mathcal{H} . Intuitively, this requires taking the time-derivative of the moments, which yields a differential equation system governing the evolution of the moment sequence, given by

$$\frac{d}{dt} m_k(t) = \frac{d}{dt} \langle \Phi_k, x_t \rangle = \left\langle \Phi_k, \frac{d}{dt} x_t \right\rangle = \langle \Phi_k, F(t, \cdot, x_t, u(t)) \rangle. \quad (7)$$

Here, the change of the order of the time-derivative and the inner product operation follows from the dominated convergence theorem (Folland, 2013). Note that because the state-space \mathcal{H} of the parameterized system is a vector space, the vector field $F(t, \beta, x_t(\beta), u(t))$ governing the system dynamics, considered as a function in β , is an element of \mathcal{H} as well. Following the definition in (6), we observe that $\langle \Phi_k, F(t, \cdot, x_t, u(t)) \rangle$ in (7) is essentially the k^{th} moment of $F(t, \cdot, x_t, u(t))$. Let $\bar{F}(t, m(t), u(t))$ denote the moment sequence of $F(t, \cdot, x_t, u(t)) = F(t, \cdot, \mathcal{K}^{-1}m(t), u(t))$. We obtain a concrete representation of the moment kernelized ensemble system defined on \mathcal{M} as

$$\frac{d}{dt} m(t) = \bar{F}(t, m(t), u(t)). \quad (8)$$

Remark 6 *Note that the moment kernelized system in (8) consists of countably many components, even though the parameterized ensemble system in (4) may be composed of a continuum of (uncountably many) intelligent agents. This implies that the moment kernelization not only defines a kernel parameterization but also provides model reduction for parameterized systems. This feature will be fully exploited in the development of the filtrated RL architecture for learning optimal policies of parameterized systems.*

Moment kernelization of value functions. The final step in exploring the moment kernel representation of the proposed policy learning problem is to kernelize the state-value function $V : [0, T] \times \mathcal{H} \rightarrow \mathbb{R}$. To achieve this, supported by the integrability condition in Assumption C1, we apply Fubini’s theorem to change the order of the integrals in the first summand of V , resulting in $V(t, x_t) = \int_t^T \int_{\Omega} r(s, \beta, x(s, \beta), u(s)) d\beta ds + \int_{\Omega} K(T, \beta, x(T, \beta)) d\beta$ (Folland, 2013). Observe that $\int_{\Omega} r(t, \beta, x(t, \beta), u(t)) d\beta$ and $\int_{\Omega} K(T, \beta, x(T, \beta)) d\beta$ are essentially the 0th moments of $r(t, \beta, x_t, u(t))$ and $K(T, \beta, x_T)$ as real-valued functions defined on Ω , provided that Φ_0 is a constant function. Denoting these by $\bar{r}(t, m(t), u(t))$ and $\bar{K}(T, m(T))$, respectively, we obtain the desired moment kernelized state-value function as

$$V(t, m(t)) = \int_t^T \bar{r}(s, m(s), u(s)) ds + \bar{K}(T, m(T)). \quad (9)$$

Remark 7 *The key step in kernelizing the state-value function above is changing the order of the integrals with respect to $\beta \in \Omega$ and $t \in [0, T]$ using the integrability condition in Assumption C1. This derivation remains valid under a relatively weaker condition, namely, that the immediate reward function r is nonnegative for any admissible control policy, as a consequence of Tonelli’s theorem (Folland, 2013). This nonnegative condition also occurs commonly in practice, e.g., the LQR problem presented in Section 2.1, demonstrating the general applicability of the proposed moment kernelization approach.*

As the pointwise infimum of the state-value function over the space of admissible policies, the value function naturally admits the moment kernel representation $V^* : [0, T] \times \mathcal{M} \rightarrow \mathbb{R}$, given by

$$V^*(t, m(t)) = \inf_{u \in \mathcal{U}} V(t, m(t)) = \inf_{u \in \mathcal{U}} \left\{ \int_t^{t+h} \bar{r}(s, m(s), u(s)) ds + V^*(t+h, m(t+h)) \right\} \quad (10)$$

for any $h > 0$ such that $t+h \leq T$, where the second equality follows from the dynamic programming principle (Evans, 2010). The integral representation of the value function in (10) leads to the following regularity property, which plays a crucial role in establishing RL approaches to policy learning of parameterized systems over the moment domain.

Proposition 8 *The value function $V^* : [0, T] \times \mathcal{M} \rightarrow \mathbb{R}$ is Lipschitz continuous.*

Proof See Appendix B. ■

3.2 Moment convergence for reinforcement learning of kernelized ensemble systems

As the moment kernelized system in (8) consists of countable state variables, this enables the use of truncated moment systems to facilitate RL of parameterized systems. This section is dedicated to laying the theoretical foundation for this approach. In particular, the main focus is to show that the value functions of truncated moment systems converge to those of the moment kernelized parameterized systems in an appropriate sense.

To rigorously formulate the corresponding policy learning problem for a truncated moment system, we use the hat notation ‘ $\hat{\cdot}$ ’ to denote the truncation operation and identify the order- N truncated moment sequence with the projection of the infinite moment sequence onto the first N components, e.g., $\hat{m}_N(t) = (m_0(t), \dots, m_N(t), 0, 0, \dots)'$ and $\hat{F}_N = (\bar{F}_0, \dots, \bar{F}_N, 0, 0, \dots)'$. This then constructs the order- N truncated moment system as

$$\frac{d}{dt}\hat{m}_N(t) = \hat{F}_N(t, \hat{m}_N(t), u(t)), \quad (11)$$

which is a control system defined on $\widehat{\mathcal{M}}_N = \{\hat{z}_N \in \mathcal{M} \mid z \in \mathcal{M}\}$, the space of all order- N truncated moment sequences. With a slight abuse of the hat notation, we denote the state-value and value functions of the system in (11) by $\widehat{V}_N[0, T] \times \widehat{\mathcal{M}}_N \rightarrow \mathbb{R}$ and $\widehat{V}_N^* : [0, T] \times \widehat{\mathcal{M}}_N \rightarrow \mathbb{R}$, respectively. Clearly, the dynamic programming principle is also satisfied by this “truncated” value function as

$$\widehat{V}_N^*(t, \hat{m}_N(t)) = \inf_{u \in \mathcal{U}} \widehat{V}_N(t, \hat{m}_N(t)) = \inf_{u \in \mathcal{U}} \int_t^T \bar{r}(s, \hat{m}_N(s), u(s)) ds + \bar{K}(T, \hat{m}_N(T)). \quad (12)$$

The RKHS structure on the moment space \mathcal{M} implies the convergence of the truncated moment sequence $\hat{m}_N(t)$ to the entire moment sequence $m(t)$ as $N \rightarrow \infty$. This can be observed by the vanishing of the truncation error as

$$\|m(t) - \hat{m}_N(t)\|^2 = \sum_{k=N}^{\infty} m_k^2(t) \rightarrow 0 \quad (13)$$

as $N \rightarrow \infty$, because $m_k \rightarrow 0$ as $k \rightarrow \infty$ following from $m(t)$ being square summable. We then have

$$\begin{aligned} \|\bar{F}(t, m(t), u(t)) - \hat{F}_N(t, \hat{m}_N(t), u(t))\| &\leq \|\bar{F}(t, m(t), u(t)) - \hat{F}_N(t, m(t), u(t))\| \\ &\quad + \|\hat{F}_N(t, m(t), u(t)) - \hat{F}_N(t, \hat{m}_N(t), u(t))\| \rightarrow 0. \end{aligned}$$

Here, because \bar{F} is also a moment sequence as discussed previously, the convergence of the first term to 0 is essentially the same as (13). Combined with the continuity of \hat{F}_N , this ensures the convergence of the second term to 0.

To reinforce the convergence property from truncated moment sequences to “truncated” value functions, it is crucial to be aware of the fact that \widehat{V}_N^* is barely the restriction of \widehat{V}^* on the subspace $\widehat{\mathcal{M}}_N \subset \mathcal{M}$, which we denote by $V^*|_{\widehat{\mathcal{M}}_N}$. This is because the state trajectory of the order- N truncated moment system in (11) must stay within $\widehat{\mathcal{M}}_N$, while the optimal trajectory of the entire moment system (8), starting from an initial condition in $\widehat{\mathcal{M}}_N$, may leave $\widehat{\mathcal{M}}_N$. Therefore, it is necessary that $\widehat{V}_N^*(t, \hat{m}_N(t)) \geq V^*|_{\widehat{\mathcal{M}}_N}(t, \hat{m}_N(t))$, but this minor discrepancy will not preclude the desired convergence of \widehat{V}_N^* to V^* as $N \rightarrow \infty$. To make this convergence argument mathematically rigorous, we extend the domain of \widehat{V}_N^* from $[0, T] \times \widehat{\mathcal{M}}_N$ to $[0, T] \times \mathcal{M}$ by defining $\widehat{V}_N^*(t, m(t)) \doteq \widehat{V}_N(t, \hat{m}_N(t))$.

Theorem 9 (Moment convergence of value functions) *The value function $(\widehat{V}_N^*)_{N \in \mathbb{N}}$ of the truncated moment system in (11) converges locally uniformly to V^* , the value function of the moment system in (8), on $[0, T] \times \mathcal{M}$ as $N \rightarrow \infty$.*

Proof Because the spaces of truncated moment sequences form an ascending chain of subspaces of the moment space, meaning $\widehat{\mathcal{M}}_0 \subset \widehat{\mathcal{M}}_1 \subset \dots \subset \mathcal{M}$, the sequence of value functions $(\widehat{V}_N^*)_{N \in \mathbb{N}}$ forms a decreasing chain $\widehat{V}_0^*(t, z) \geq \widehat{V}_1^*(t, z) \geq \dots \geq V^*(t, z)$ for any $z \in \mathcal{M}$. This particularly implies that $(\widehat{V}_N^*)_{N \in \mathbb{N}}$ is locally uniformly bounded, i.e., there exists a neighborhood \mathcal{N} of (t, z) in $[0, T] \times \mathcal{M}$ and a real number M so that $|\widehat{V}_N^*(s, w)| \leq M$ for all $(s, w) \in \mathcal{N}$ and $N \in \mathbb{N}$ with M independent of (s, w) and N . On the other hand, by the aforementioned fact that $\widehat{V}_N^* \geq V^*$, together with the Lipschitz continuity of V^* shown in Lemma 8, all the functions in the sequence $(\widehat{V}_N^*)_{N \in \mathbb{N}}$ are Lipschitz continuous and have the same Lipschitz constant equal to that of V^* . Consequence, this sequence of value functions is equicontinuous Rudin (1976). A direct application of the Arzelà–Ascoli theorem then shows that $\widehat{V}_N^* \rightarrow V^*$ on \mathcal{N} uniformly as $N \rightarrow \infty$ (Folland, 2013). Since $(t, z) \in [0, T] \times \mathcal{M}$ is arbitrary, we obtain $\widehat{V}_N^* \rightarrow V^*$ locally uniformly on $[0, T] \times \mathcal{M}$ as $N \rightarrow \infty$ as desired. ■

According to the dynamic programming principle illustrated in (10), value functions are obtained by evaluating the corresponding state-value functions along optimal trajectories. It then becomes intuitive that, accompanied by the convergence of the “truncated” value function, the optimal trajectory of the truncated moment system also converges to that of the entire moment system.

Theorem 10 (Moment convergence of optimal trajectories) *Let $\widehat{m}_N^*(t)$ and $m^*(t)$ be the optimal trajectories of the order- N truncated moment system and the entire moment system in (11) and (8), respectively. Then, $\widehat{m}_N^*(t) \rightarrow m^*(t)$ on \mathcal{M} as $N \rightarrow \infty$ for all $t \in [0, T]$. Equivalently, $\widehat{x}_N^*(t, \cdot) \rightarrow x^*(t, \cdot)$ on \mathcal{H} as $N \rightarrow \infty$ for all $t \in [0, T]$, where $\widehat{x}_N^*(t, \cdot)$ and $x^*(t, \cdot)$ are the trajectories of the parameterized system in (4) driven by the optimal control policies for the truncated and entire moment systems, respectively.*

Proof By the dynamic programming principle in (10), the optimal trajectory $m^*(t)$ on $[0, T]$ remains optimal when restricted to any subinterval $[s, T]$ for $0 < s < T$, and the same result holds for $\widehat{m}_N^*(t)$ for each truncation order $N \in \mathbb{N}$. Because the sequence of real numbers $(\widehat{V}_N(t, \widehat{m}_N^*(t)))_{N \in \mathbb{N}}$ is monotonically decreasing and bounded from below by $V(t, m^*(t))$ as shown in the proof of Theorem 9, it is necessary that $\widehat{V}_N(t, \widehat{m}_N^*(t)) \rightarrow V(t, m^*(t))$ as $N \rightarrow \infty$ for each $t \in [0, T]$ (Rudin, 1976). Together with the continuity of V and the locally uniform convergence of the value function sequence $(\widehat{V}_N)_{N \in \mathbb{N}}$ by Theorem 9, we obtain $V(t, m^*(t)) = \lim_{N \rightarrow \infty} \widehat{V}_N(t, \widehat{m}_N^*(t)) = V(t, \lim_{N \rightarrow \infty} \widehat{m}_N^*(t))$ so that $\widehat{m}_N^*(t) \rightarrow m^*(t)$ on \mathcal{M} as $N \rightarrow \infty$. The isometrically isomorphic property of the moment kernel transform $\mathcal{K} : \mathcal{H} \rightarrow \mathcal{M}$ then implies that $\mathcal{K}^{-1}\widehat{m}_N^*(t) \rightarrow \mathcal{K}^{-1}m^*(t)$, i.e., $\widehat{x}_N^*(t, \cdot) \rightarrow x^*(t, \cdot)$, on \mathcal{H} as $N \rightarrow \infty$. ■

A fundamental property of value functions crucial to RL is that they are viscosity solutions of Hamiltonian-Jacobi-Bellman (HJB) equations (Evans, 2010), which enables the design of RL algorithms to learn optimal control policies. The moment convergence shown in Theorems 9 and 10 then leads to an extension of HJB equations to value functions defined on infinite-dimensional spaces.

Corollary 11 (Moment convergence of Hamilton-Jacobi-Bellman equations)

The value function $V^* : [0, T] \times \mathcal{M} \rightarrow \mathbb{R}$ is the unique viscosity solution of the Hamilton-Jacobi-Bellman equation, given by

$$\frac{\partial}{\partial t} V^*(t, z) + \min_{a \in U} \left\{ \langle DV^*(t, z), \bar{F}(t, z, a) \rangle + \bar{r}(t, z, a) \right\} = 0, \quad (14)$$

on $(0, T) \times \mathcal{M}$ with the boundary condition $V = K$ on $\{t = T\} \times \mathcal{M}$, where $DV^*(t, z)$ is the (Gateaux) differential of V^* with respect to $z \in \mathcal{M}$.

Proof The uniqueness directly follows from the Lipschitz continuity of the Hamiltonian $H(t, z, p) = \min_{a \in U} \{ \langle p, \bar{F}(t, z, a) \rangle + \bar{r}(t, z, a) \}$ in $(z, p) \in \mathcal{M} \times \mathcal{M}^*$ uniformly in $t \in [0, T]$, as a consequence of Assumption C2, where \mathcal{M}^* is the dual space of \mathcal{M} (Evans, 2010). Then, it remains to examine whether V^* is a viscosity solution of the first-order partial differential equation in (14).

To show this, for any $(t, z) \in [0, T] \times \mathcal{M}$, we pick a continuously differentiable function $v : [0, T] \times \mathcal{M} \rightarrow \mathbb{R}$ such that $V - v$ has a local maximum at (t, z) . Without loss of generality, we can assume (t, z) to be a strict local maximum; that is, there exists a neighborhood \mathcal{N} whose closure $\bar{\mathcal{N}}$ contains (t, z) such that $V^*(t, z) - v(t, z) > V^*(s, w) - v(t, z)$ for any $(s, w) \in \bar{\mathcal{N}} \setminus \{(t, z)\}$. Let (t_N, z_N) be a maximum of $\hat{V}_N^* - v$ on $\bar{\mathcal{N}}$. Then, we have

$$\frac{\partial}{\partial t} v(t_N, z_N) + \min_{a \in U} \left\{ \langle Dv(t_N, z_N), \hat{F}_N(t_N, z_N, a) \rangle + \bar{r}(t_N, z_N, a) \right\} \geq 0 \quad (15)$$

following from the fact that \hat{V}_N^* is the value function for the truncated moment system in (11) defined on the finite-dimensional space $\hat{\mathcal{M}}_N$ (Evans, 2010).

By passing to a subsequence and shrinking the neighborhood \mathcal{N} if necessary, we have $(t_N, z_N) \rightarrow (\tilde{t}, \tilde{z})$ as $N \rightarrow \infty$ for some $\tilde{t} \in [0, T]$ and $\tilde{z} \in \bar{\mathcal{N}}$. This leads to the convergence $\hat{V}_N^*(t_N, z_N) - v(t_N, z_N) \rightarrow V^*(\tilde{t}, \tilde{z}) - v(\tilde{t}, \tilde{z})$ by the locally uniform convergence of \hat{V}_N^* to V^* shown in Theorem 9. Because $\hat{V}_N^*(t_N, z_N) - v(t_N, z_N) \geq \hat{V}_N^*(s, w) - v(s, w)$ for any $(s, w) \in \bar{\mathcal{N}}$ and $N \in \mathbb{N}$ by the choice of (t_N, z_N) , we obtain $V^*(\tilde{t}, \tilde{z}) - v(\tilde{t}, \tilde{z}) \geq V^*(s, w) - v(s, w)$ for any $(s, w) \in \bar{\mathcal{N}}$ by letting $N \rightarrow \infty$. In particular, $V^*(\tilde{t}, \tilde{z}) - v(\tilde{t}, \tilde{z}) \geq V^*(t, z) - v(t, z)$. This shows that $(\tilde{t}, \tilde{z}) = (t, z)$ and hence $(t_N, z_N) \rightarrow (t, z)$ as $N \rightarrow \infty$.

Passing to the limit as $N \rightarrow \infty$ in (15) yields

$$\frac{\partial}{\partial t} v(t, z) + \min_{a \in U} \left\{ \langle Dv(t, z), \bar{F}(t, z, a) \rangle + \bar{r}(t, z, a) \right\} \geq 0, \quad (16)$$

where we use the uniform convergence of \hat{F}_N to \bar{F} , following from a similar proof as Theorem 9 by replacing \hat{V}_N^* and V^* with \hat{F}_N and \bar{F} , respectively, and the continuity of the Hamiltonian $H(t, z, p) = \min_{a \in U} \{ \langle p, \bar{F}(t, z, a) \rangle + \bar{r}(t, z, a) \}$. A similar argument shows that

$$\frac{\partial}{\partial t} w(t, z) + \min_{a \in U} \left\{ \langle Dw(t, z), \bar{F}(t, z, a) \rangle + \bar{r}(t, z, a) \right\} \leq 0 \quad (17)$$

if $V^* - w$ attains a local minimum at $(t, z) \in [0, T] \times \mathcal{M}$. The equations in (16) and (17) imply that V^* is a viscosity solution of the Hamilton-Jacobi-Bellman equation in (14). ■

3.3 Moment Kernelization of Parameterized Systems in Stochastic Environments

The proposed parametrized model formulation and moment kernelization framework naturally extend to stochastic settings. We explore the adoption of our approach in three types of stochastic environments.

Model uncertainty. The parametrized ensemble system in (4) also provides a natural formulation for describing systems with model uncertainties. In this interpretation, the system parameter β characterizes uncertainty and can be regarded as a random variable drawn from a probability distribution \mathbb{P} on the parameter space Ω . Consequently, the trajectory of the parameterized system x_t is a stochastic process on the probability space (Ω, \mathbb{P}) , and the value function becomes

$$V(t, x_t) = \mathbb{E} \left(\int_t^T r(s, x_s, u(s)) ds + K(T, x_T) \right) = \int_{\Omega} \left(\int_t^T r(s, x_s, u(s)) ds + K(T, x_T) \right) d\mathbb{P},$$

where \mathbb{E} denotes the expectation with respect to the probability measure \mathbb{P} . Suppose that x_t possesses finite variance for all t , then the k^{th} moment of the parameterized system can be defined as

$$m_k(t) = \langle \Phi_k, x_t \rangle = \mathbb{E}(\Phi_k x_t) = \int_{\Omega} \Phi_k x_t d\mathbb{P}, \quad (18)$$

where $\{\Phi_k\}_{k \in \mathbb{N}}$ is an orthonormal basis of $L^2(\Omega, \mathbb{P})$. Leveraging this definition of moments, the kernelized system and value function follow the same forms as in (8) and (9), respectively. This implies that all the developments and conclusions in Sections 3.1 and 3.2 carry over to the stochastic setting with model uncertainties.

Background noise. Background noise is frequently inherent in the agents' environments and typically appears as additive noise present in measurements of the states of agents. Here, we consider two types of background noise commonly encountered in practice: independent noise and common noise.

- Independent noise: In this case, each agent in the parameterized family in (4) experiences parameter-dependent independent additive noise of the form $x_t(\beta) + \varepsilon(\beta)$, where $\{\varepsilon(\beta)\}_{\beta \in \Omega}$ is a family of pairwise independent random variables. Suppose that $\varepsilon(\beta)$ has zero mean and finite variance for all $\beta \in \Omega$. Then, Fubini's theorem (Folland, 2013) implies that

$$\begin{aligned} \mathbb{E} \left(\int_{\Omega} \varepsilon(\beta) d\beta \right)^2 &= \mathbb{E} \left(\int_{\Omega} \varepsilon(\beta) d\beta \int_{\Omega} \varepsilon(\gamma) d\gamma \right) = \mathbb{E} \left(\int_{\Omega^2} \varepsilon(\beta) \varepsilon(\gamma) d\beta d\gamma \right) \\ &= \int_{\Omega^2} \mathbb{E}(\varepsilon(\beta) \varepsilon(\gamma)) d\beta d\gamma = \int_D \mathbb{E}(\varepsilon^2(\beta)) d\beta d\gamma + \int_{\Omega^2 \setminus D} \mathbb{E}(\varepsilon(\beta)) \mathbb{E}(\varepsilon(\gamma)) d\beta d\gamma. \end{aligned}$$

Here, $D = \{(\beta, \gamma) \in \Omega^2 : \beta = \gamma\}$ is the diagonal subset of $\Omega^2 = \Omega \times \Omega$, and the last equality follows from the pairwise independence of $\{\varepsilon(\beta)\}_{\beta \in \Omega}$. We further observe that D has Lebesgue measure 0, and together with $\mathbb{E}(\varepsilon^2(\beta)) < \infty$, this leads to $\int_D \mathbb{E}(\varepsilon^2(\beta)) d\beta d\gamma = 0$. On the other hand, the zero mean property gives

$\int_{\Omega^2 \setminus D} \mathbb{E}(\varepsilon(\beta)) \mathbb{E}(\varepsilon(\gamma)) d\beta d\gamma = 0$. As a result, we have $\mathbb{E} \left(\int_{\Omega} \varepsilon(\beta) d\beta \right)^2 = 0$, indicating that $\int_{\Omega} \varepsilon(\beta) d\beta = 0$ almost surely (*a.s.*).

Following the same definition as in (6), the k^{th} moment satisfies

$$m_k(t) = \langle \Phi_k, x_t + \varepsilon \rangle = \int_{\Omega} \Phi_k(\beta) x_t(\beta) d\beta + \int_{\Omega} \Phi_k(\beta) \varepsilon(\beta) d\beta = \langle \Phi_k, x_t \rangle \quad a.s.$$

for all $k \in \mathbb{N}$. Here, we use the claim that $\int_{\Omega} \Phi_k(\beta) \varepsilon(\beta) d\beta = 0$ *a.s.* To see this, consider that because of the square-integrability of Φ_k , the set $L = \{\beta \in \Omega : \Phi_k(\beta) = \infty\}$ has Lebesgue measure 0. This gives $\int_{\Omega} \Phi_k(\beta) \varepsilon(\beta) d\beta = \int_L \Phi_k(\beta) \varepsilon(\beta) d\beta + \int_{\Omega \setminus L} \Phi_k(\beta) \cdot 0 d\beta = 0$ *a.s.*, as claimed.

- **Common noise:** In this case, each agent in the parameterized ensemble system in (4) experiences a common noise ε , so that its environmental state follows the form $x_t(\beta) + \varepsilon$. Suppose that ε is a random variable with zero mean and finite variance, then the k^{th} moment is given by

$$m_k(t) = \langle \Phi_k, x_t + \varepsilon \rangle = \int_{\Omega} \Phi_k(\beta) x_t(\beta) d\beta + \varepsilon \int_{\Omega} \Phi_k(\beta) d\beta,$$

is a random variable with the mean given by

$$\mathbb{E}(m_k(t)) = \int_{\Omega} \Phi_k(\beta) x_t(\beta) d\beta + \int_{\Omega} \Phi_k(\beta) d\beta \cdot \mathbb{E}(\varepsilon) = \int_{\Omega} \Phi_k(\beta) x_t(\beta) d\beta = \langle \Phi_k, x_t \rangle.$$

Statistically, this implies that the moments of the parameterized system in the presence of common background noise are an unbiased estimator of the moment without noise. More importantly, the moment kernelization reduces the variance. By Hölder's inequality, we have

$$\text{Var}(m_k(t)) = \left(\int_{\Omega} \Phi_k(\beta) d\beta \right)^2 \text{Var}(\varepsilon) \leq |\Omega| \|\Phi_k\|^2 \text{Var}(\varepsilon) < \text{Var}(\varepsilon),$$

provided, without loss of generality, the Lebesgue measure of Ω satisfies $|\Omega| < 1$.

Stochastic dynamics. When the parameterized system in (4) is driven by a noise process, e.g., a Brownian motion W_t on a probability space $(\mathbb{R}^n, \mathbb{P})$, the evolution of its state obeys the stochastic differential equation

$$dx_t(\beta) = F(t, \beta, x_t(\beta), u(t)) dt + G(t, \beta, x_t(\beta)) dW_t,$$

with the state-value function given by

$$V(t, x_t) = \mathbb{E} \left(\int_{\Omega} \left[\int_t^T r(s, x_s, u(s)) ds + K(T, x_T) \right] d\beta \right),$$

where F and G satisfy the Lipschitz and linear growth conditions in the system state variable, $\|F(t, \beta, x, u) - F(t, \beta, y, u)\| + \|G(t, \beta, x, u) - G(t, \beta, y, u)\| \leq K_1 \|x - y\|$ and $\|F(t, \beta, x, u)\|^2 + \|G(t, \beta, x, u)\|^2 \leq K_2(1 + \|x\|^2)$, respectively, uniformly in $(t, \beta, u) \in [0, T] \times \Omega \times \mathcal{U}$. These

conditions guarantee that, for each $\beta \in \Omega$, the system trajectory is a finite variance stochastic process on the probability space $(\mathbb{R}^n, \mathbb{P})$ (Øksendal, 2003). Following the definition in (18), the moments of this parameterized stochastic system can be defined as

$$m_{kl}(t) = \mathbb{E}(\Phi_k \Psi_l \circ x_t) = \mathbb{E}\left(\int_{\Omega} \Phi_k(\beta) \Psi_l(x_t(\beta)) d\beta\right)$$

for all $(k, l) \in \mathbb{N}^2$, where $(\Psi_l)_{l \in \mathbb{N}}$ is an orthonormal basis of $L^2(\mathbb{R}^n, \mathbb{P})$. Note that, in this case, the moment sequence $(m_{kl}(t))_{k, l \in \mathbb{N}}$ is a deterministic double sequence, demonstrating again the ability of the moment kernel transform to mitigate the noise effect.

It should be noted that, within a stochastic environment, the open-loop requirement for the control policy may cause the loss of the Markov property in the state evolution of the parameterized system. This severely limits the applicability of MDP-based RL algorithms to the parameterized system, underscoring the necessity for the proposed moment kernelization technique to reduce randomness involved in the system. In the next section, we will develop a new RL architecture that fully exploits the algebraic structure of moment kernelized systems to facilitate policy learning for parameterized systems.

4. Filtrated Reinforcement Learning Architecture for Policy Learning of Parameterized Ensemble Systems

Building upon the theoretical foundations established in the previous sections, we now turn our attention to algorithmic approaches to policy learning in parameterized ensemble systems evolving on infinite-dimensional function spaces. In this section, we will develop a novel RL architecture with effective algorithms by organizing truncated moment kernelized systems with increasing truncation orders into a filtrated structure. We will adopt spectral sequence techniques to conduct a convergence analysis of the proposed filtrated RL (FRL) algorithms. Additionally, we will quantitatively investigate the computational and sample efficiency of FRL. The performance and efficiency of the FRL algorithms will be demonstrated using examples arising from practical applications and compared with baseline deep RL models.

4.1 Filtrated policy search for moment kernelized systems

After transforming the parameterized ensemble system to the moment domain, it is evident that an order- N truncated moment system contains any truncated moment system with a lower truncation order $N' < N$ as a subsystem. From a learning perspective, the corresponding RL problem for the order- N' truncated moment system is a subproblem of that for the order- N truncated moment system. Generally, along any increasing sequence of truncation orders $N_0 < N_1 < \dots$, a filtration consisting of an ascending chain of RL problems is revealed. To further elaborate on this FRL architecture, starting from an initial moment truncation order N_0 , an optimal policy $u_{N_0}^*(t)$ and the value function $\hat{V}_{N_0}^*(t, \hat{m}_{N_0}(t))$ can be learned for the order- N_0 truncated moment system. Next, the truncation order is increased to N_1 to learn $u_{N_1}^*(t)$ and $\hat{V}_{N_1}^*(t, \hat{m}_{N_1}(t))$ for the order- N_1 truncated moment system. Continuing this procedure, we generate a sequence of control policy $(u_{N_k}^*(t))_{k \in \mathbb{N}}$ and value function $(\hat{V}^*(t, \hat{m}_{N_k}(t)))_{k \in \mathbb{N}}$ pairs. Guaranteed by the moment convergence proved in Theorem 10, the value function sequence $(\hat{V}^*(t, \hat{m}_{N_k}(t)))_{k \in \mathbb{N}}$ necessarily converges to the

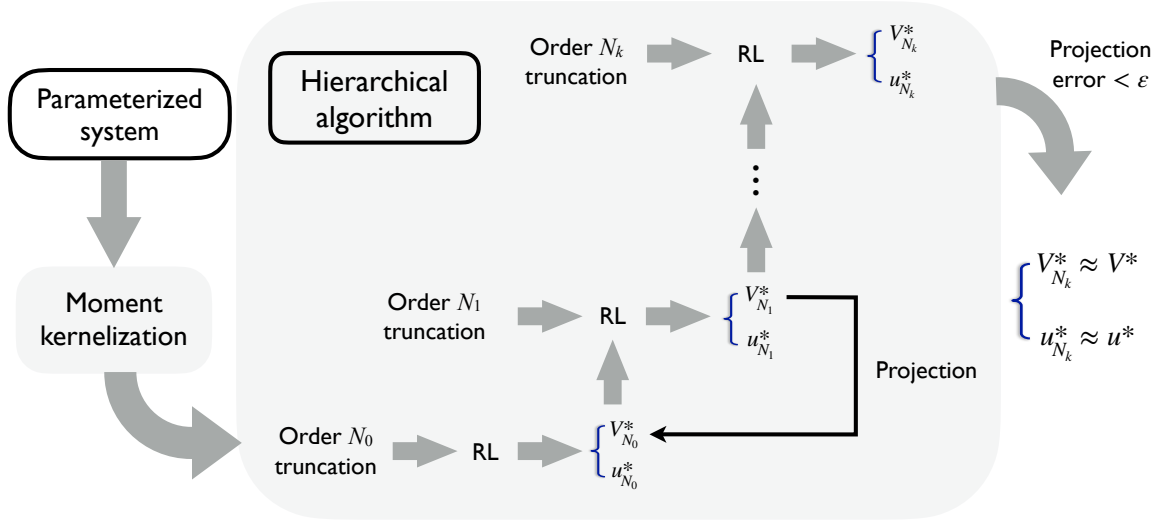


Figure 2: Workflow of FRL for parameterized ensemble systems defined on infinite-dimensional function spaces.

value function $V^*(t, m(t))$ of the entire moment system. Algorithmically, when the *projection error* $\sup_{t \in [0, T]} |\hat{V}_{N_k}^*(t, \hat{m}_{N_k}^*(t)) - \hat{V}_{N_{k-1}}^*(t, \hat{m}_{N_{k-1}}^*(t))| < \epsilon$ is satisfied for a prescribed tolerance $\epsilon > 0$, the optimal policy $u_{N_k}^*(t)$ of the order- N_k truncated moment system provides a sufficiently good approximation to the optimal policy $u^*(t)$ of the entire moment system, equivalently the parameterized ensemble system. The workflow of this FRL approach is shown in Figure 2.

Remark 12 (Optimality preserving filtration) *The observation made above Theorem 9, namely, $\hat{V}_N^* \geq V^*|_{\widehat{M}_N}$ for any moment truncation order $N \in \mathbb{N}$, directly extends to the established RL filtration; specifically $\hat{V}_{N_i}^* \geq \hat{V}_{N_j}^*|_{\widehat{M}_{N_i}} \geq \hat{V}_{N_j}^*$ for any $i < j$. This means that the optimal policy learned at one hierarchy of the filtrated RL problem is also optimal for, and may even decrease the values of, the cumulative reward function learned from all lower-level hierarchies. This demonstrates that moving up the hierarchy in the RL filtration preserves the optimality of each preceding hierarchy.*

A direct observation of the truncated and entire moment systems in (11) and (8) reveals that they are regulated by the same control policy. This highlights policy search algorithms as prime candidates for learning the optimal policy at each hierarchy in the RL filtration. Moreover, the optimality-preserving property revealed in Remark 12 strongly suggests a specific algorithmic approach: the policy learned from the current hierarchy is always guaranteed to be a good initial condition for the successive hierarchy. Consequently, as the truncation order increases, the initial condition becomes closer to the optimal solution, effectively reducing the computational cost for learning optimal policies of high-dimensional systems. This filtrated policy search algorithm is summarized in Algorithm 1.

Algorithm 1 Filtrated policy search for learning optimal policies for parameterized systems

Input: Initial state x_0 , final time T , projection error tolerance ε

Output: Optimal policy u^*

Initialization: truncation order N_0 , initial policy $u^{(0)}$, projection error $P > \varepsilon$, hierarchy level $i = 0$

```

2: while  $P > \varepsilon$  do
    if  $i > 0$  then
4:          $i \leftarrow i + 1$ 
        Pick  $N_i > N_{i-1}$  and  $u^{(0)} \leftarrow u_{N_{i-1}}^*$ 
6:     end if
        Generate data by solving the parameterized system with the input  $u^{(0)}$  and collect the rewards
8:     Compute order  $N_i$  truncated moment kernelization:  $\hat{m}_{N_i}(t)$ ,  $\hat{F}_{N_i}(t, \hat{m}_{N_i}(t), u^{(0)})$ ,  $\bar{r}(t, \hat{m}_{N_i}(t), u^{(0)})$ ,  $\bar{K}(T, \hat{m}_{N_i}(T))$ 
        Solve  $u_{N_i}^*(t) = \arg\min_u \hat{V}_{N_i}(t, \hat{m}_{N_i}(t))$  and compute  $\hat{V}_{N_i}^*(t, \hat{m}_{N_i}^*(t)) = \min_u \hat{V}_{N_i}^*(t, \hat{m}_{N_i}(t))$ 
10:    Compute  $P = \sup_{t \in [0, T]} |\hat{V}_{N_i}^*(t, \hat{m}_{N_i}^*(t)) - \hat{V}_{N_{i-1}}^*(t, \hat{m}_{N_{i-1}}^*(t))|$  with  $\hat{V}_{N_{-1}}^* \equiv 0$ 
    end while
12:  $u^* \leftarrow u_{N_i}^*$  return  $u^*$ 
    
```

4.1.1 SPECTRAL SEQUENCE CONVERGENCE OF FILTRATED POLICY SEARCH

Let $\hat{m}_{N_i}^{(j)}(t)$ denote the state trajectory of the order- N_i truncated moment system driven by the policy $u_{N_i}^{(j)}$ resulting from the j^{th} iteration of the policy search (PS) algorithm. The filtrated policy search shown in Algorithm 1 then generates a *spectral sequence*, given by

	PS iteration 0	PS iteration 1	...	PS iteration j	...
Hierarchy 0	$\hat{V}_{N_0}(t, \hat{m}_{N_0}^{(0)}(t))$	$\hat{V}_{N_0}(t, \hat{m}_{N_0}^{(1)}(t))$...	$\hat{V}_{N_0}(t, \hat{m}_{N_0}^{(j)}(t))$	$\longrightarrow \hat{V}_{N_0}^*(t, \hat{m}_{N_0}(t))$
Hierarchy 1	$\hat{V}_{N_1}(t, \hat{m}_{N_1}^{(0)}(t))$	$\hat{V}_{N_1}(t, \hat{m}_{N_1}^{(1)}(t))$...	$\hat{V}_{N_1}(t, \hat{m}_{N_1}^{(j)}(t))$	$\longrightarrow \hat{V}_{N_1}^*(t, \hat{m}_{N_1}(t))$
\vdots	\vdots	\vdots		\vdots	\vdots
Hierarchy i	$\hat{V}_{N_i}(t, \hat{m}_{N_i}^{(0)}(t))$	$\hat{V}_{N_i}(t, \hat{m}_{N_i}^{(1)}(t))$...	$\hat{V}_{N_i}(t, \hat{m}_{N_i}^{(j)}(t))$	$\longrightarrow \hat{V}_{N_i}^*(t, \hat{m}_{N_i}(t))$
\vdots	\downarrow	\downarrow		\downarrow	\downarrow
	$V(t, m^{(0)}(t))$	$V(t, m^{(1)}(t))$...	$V(t, m^{(j)}(t))$	$\longrightarrow V^*(t, m(t))$

In this sequence, the i^{th} row converges to the value function $\hat{V}_{N_i}^*(t, \hat{m}_{N_i}(t))$ of order- N_i truncated moment system, and the j^{th} column converges to the state-value function $V(t, m^{(j)}(t))$ of the entire moment system. Specifically, the row convergence naturally follows from the convergence of the PS algorithm (Sutton and Barto, 2018), while the column convergence results from the continuity of the state-value function and the convergence of the truncated moment sequence to the infinite moment sequence. The row and column convergence together imply a stronger convergence property of Algorithm 1, that is, *spectral sequence convergence*, as defined and proved below.

Theorem 13 (Spectral sequence convergence of FRL) *Let $(\widehat{V}_{N_i}(t, \widehat{m}_{N_i}^{(j)}))_{i,j \in \mathbb{N}}$ be the spectral sequence generated by Algorithm 1. Then, this sequence converges to the value function $V^*(t, m(t))$ of the moment kernelized system, i.e., $\widehat{V}_i(t, \widehat{m}_{N_i}^{(\phi(i))}) \rightarrow V^*(t, m(t))$ as $i \rightarrow \infty$ for any monotonically increasing function $\phi : \mathbb{N} \rightarrow \mathbb{N}$.*

Proof It suffices to show that both iterated limits, $\lim_{i \rightarrow \infty} \lim_{j \rightarrow \infty} \widehat{V}_{N_i}(t, \widehat{m}_{N_i}^{(j)})$ and $\lim_{j \rightarrow \infty} \lim_{i \rightarrow \infty} \widehat{V}_{N_i}(t, \widehat{m}_{N_i}^{(j)})$, exist and are equal to $V^*(t, m(t))$.

We first compute $\lim_{i \rightarrow \infty} \lim_{j \rightarrow \infty} \widehat{V}_{N_i}(t, \widehat{m}_{N_i}^{(j)})$. For each $j \in \mathbb{N}$, following a similar proof as for Theorem 10 by replacing $\widehat{m}_N^*(t)$ and $m^*(t)$ with $\widehat{m}_{N_i}^{(j)}(t)$ and $m^{(j)}(t)$, respectively, we obtain $\widehat{V}_{N_i}(t, \widehat{m}_{N_i}^{(j)}) \rightarrow V(t, m^{(j)}(t))$. Next, the convergence of $V(t, m^{(j)}(t))$ to $V^*(t, m(t))$ is essentially the convergence of the policy search algorithm (Sutton and Barto, 2018), concluding that $\lim_{i \rightarrow \infty} \lim_{j \rightarrow \infty} \widehat{V}_{N_i}(t, \widehat{m}_{N_i}^{(j)}) = V^*(t, m(t))$.

On the other hand, to compute $\lim_{j \rightarrow \infty} \lim_{i \rightarrow \infty} \widehat{V}_{N_i}(t, \widehat{m}_{N_i}^{(j)})$, we first note that for each i , the convergence of $\widehat{V}_{N_i}(t, \widehat{m}_{N_i}(t))^{(j)}$ to $\widehat{V}_{N_i}^*(t, \widehat{m}_{N_i}(t))$ as $j \rightarrow \infty$ again follows from the convergence of the policy search algorithm. Since each $V_{N_i}^*(t, \widehat{m}_{N_i}(t))$, as a value function, is a viscosity solution of the HJB equation associated with the order- N_i truncated moment system, Corollary 11 implies that $V_{N_i}^*(t, \widehat{m}_{N_i}(t)) \rightarrow V^*(t, m(t))$ as $i \rightarrow \infty$. This shows that $\lim_{j \rightarrow \infty} \lim_{i \rightarrow \infty} \widehat{V}_{N_i}(t, \widehat{m}_{N_i}^{(j)}) = V^*(t, m(t))$, concluding the proof. \blacksquare

Spectral sequence-enabled early stopping for FRL. Interpreted using the spectral sequence above, Algorithm 1 approaches the optimal policy u^* of the parameterized ensemble system “along the rightmost column,” meaning the learning sequence consists of the optimal policies of all the truncated moment systems. The spectral sequence convergence of FRL shown in Theorem 13 guarantees that learning sequences along any paths towards the bottom right element V^* all converge to u^* . Algorithmically, such a learning sequence is obtained by employing early stopping in the PS algorithm at each hierarchy. For example, in the extreme case, the learning sequence along the diagonal of the spectral sequence is generated by executing only one iteration of the PS algorithm at each hierarchy. However, this “naive” diagonal learning sequence is generally not effective in terms of computational efficiency and learning accuracy. This motivates the exploration of a sophisticated early stopping criterion to fully exploit the advantage of the distinguished spectral sequence convergence of FRL.

To collect some thoughts about the design of early stopping criteria, we devote our attention to policy gradient (PG) methods, which are commonly used policy search approaches in RL literature. In the context of the spectral sequence, a PG algorithm applied at the i^{th} hierarchy generates a policy sequence in the form of $u_{N_i}^{(j)} = u_{N_i}^{(j-1)} + \delta u_{N_i}^{(j-1)}$. The update rule $\delta u_{N_i}^{(j-1)}$ is generally chosen to be proportional to $D\widehat{J}_{N_i}(u_{N_i}^{(j-1)})$, the gradient of the total reward function of the order- N_i truncated moment system evaluated at $u_{N_i}^{(j-1)}$, to ensure $u_{N_i}^{(j)} \rightarrow u_{N_i}^*$. In practice, to keep the agents’ behavior under control, e.g., staying in the safety region, the PG algorithm is generally implemented in a “clipped” manner to bound the amplitude of $\delta u_{N_i}^{(j)}$, such as Proximal Policy Optimization (PPO) and Trust Region Policy Optimization (TRPO) (Schulman et al., 2017, 2015). Inspired by this, we choose the early stopping criterion, for

moving to the successive hierarchy of FRL, to be a threshold ($\delta > 0$) for the variation of the state-value function, such that $\sup_{t \in [0, T]} |\widehat{V}_{N_k}(t, \widehat{m}_k^{(i)}(t)) - \widehat{V}_{N_k}(t, \widehat{m}_k^{(i)}(t))| > \delta$. It is worth mentioning that the use of the state-value function threshold, instead of a policy threshold as in TRPO and PPO, takes into consideration the possible failure of convergence in the generated policy sequence, as pointed out at the end of Section 2. Before integrating the hierarchy-wise early stopping criterion into the FRL algorithm, we first carry out a detailed investigation into the computational and sample efficiency of FRL.

Convergence rate. It is intuitive that the rate of convergence of the spectral sequence depends on both the row and column convergence rates. As described in the proof of Theorem 13, each row is generated by a standard policy search algorithm, and hence the rate of row convergence is entirely determined by the applied algorithm. Meanwhile, the column convergence is guaranteed by the moment convergence of the value function shown in Theorem 9, which is a consequence of the convergence of the truncated moment sequence to the entire infinite moment sequence. Hence, we first evaluate the convergence rate of the truncate moment sequence. To this end, we note that $\|\widehat{m}_N(t) - m(t)\|^2 = \sum_{k=N+1}^{\infty} |m_k|^2$ is essentially the tail of the moment sequence. Therefore, this convergence rate coincides with the rate of convergence of $m_k(t)$ to 0. Owing to the ℓ^2 -convergence of the moment sequence $\sum_{k=0}^{\infty} |m_k(t)|^2 < \infty$, by the comparison test (Rudin, 1976), it is necessary that $|m_k(t)|^2 < k^{-1}$ for large enough k since the harmonic series $\sum_{k=1}^{\infty} k^{-1}$ fails to converge. As a result, the convergence rate of $\widehat{m}_N(t)$ to $m(t)$ is bounded above by $O(N^{-1/2})$.

To leverage this to compute the convergence rate of the column convergence, we have

$$\begin{aligned}
 & |V(t, m^{(j)}(t)) - \widehat{V}_{N_i}(t, \widehat{m}_j(t))| \\
 &= \left| \int_0^T \bar{r}(s, m^{(j)}(s), u(s)) ds + \bar{K}(T, m^{(j)}(s)) - \int_0^T \bar{r}(s, \widehat{m}_{N_i}^{(j)}(s), u_{N_i}(s)) ds + \bar{K}(T, m_{N_i}^{(j)}(s)) \right| \\
 &\leq \int_0^T |r(s, m^{(j)}(s), u(s)) - \bar{r}(s, \widehat{m}_{N_i}^{(j)}(s), u_{N_i}(s))| ds + |\bar{K}(T, m^{(j)}(s)) - \bar{K}(T, m_{N_i}^{(j)}(s))| \\
 &\leq \int_0^T L_r \|m^{(j)}(s) - m_{N_i}^{(j)}(s)\| ds + L_K \|m^{(j)}(T) - m_{N_i}^{(j)}(T)\| \\
 &\leq L_r T O(N_i^{-1/2}) + L_K O(N_i^{-1/2}) \sim O(N_i^{-1/2}),
 \end{aligned}$$

where we used the Lipchitz continuity of the running cost \bar{r} and terminal cost \bar{K} , as presented in Assumption C2, with L_r and L_K denoting their Lipchitz constants, respectively.

We can now integrate the rates of the column and row convergence. Let $O(\alpha(N))$ be the convergence rate of the policy search algorithm applied to the order- N truncated moment system, then we have

$$\begin{aligned}
 & |V^*(t, m(t)) - V_N(t, \widehat{m}_N^{(\lceil 1/\alpha(N) \rceil)}(t))| \leq |V^*(t, m(t)) - V_N^*(t, \widehat{m}_N(t))| \\
 & \quad + |V_N^*(t, \widehat{m}_N(t)) - V_N(t, \widehat{m}_N^{(\lceil 1/\alpha(N) \rceil)}(t))| \sim O(N^{-1/2}) + O(\alpha(N))
 \end{aligned}$$

by the triangle inequality, where $\lceil 1/\alpha(N) \rceil$ denotes the smallest integer greater than or equal to $1/\alpha(N)$. This concludes that, in the worst-case scenario, the convergence rate of the spectral sequence is bounded above by $O(N^{-1/2}) + O(\alpha(N))$. To further elaborate on it, if the stopping criterion is set to be ε , then we should choose N satisfying $N^{-1/2} + \alpha(N) < \varepsilon$

as the moment truncation order and run $\lceil 1/\alpha(N) \rceil$ iterations of the policy search algorithm on the truncated moment system.

Computational and sample efficiency. As it is impractical to collect comprehensive measurement data from a parameterized ensemble system, practical implementations of the proposed FRL method require estimations of moments and moment systems using finite data points, even in cases where the mathematical model of the parameterized system is known. Suppose that only the states of q agents in the parameterized ensemble system in (4), say $x(t, \beta_1), \dots, x(t, \beta_q)$, can be measured. We then define the *sample moments* by

$$\mathbf{m}_k(t) = \frac{|\Omega|}{q} \sum_{i=1}^q \Phi_k(\beta_i) x(t, \beta_i) \quad (19)$$

for all $k \in \mathbb{N}$, where $|\Omega|$ is the Lebesgue measure (volume) of Ω .

Proposition 14 *Suppose that, at each time t , measurement data for q randomly selected systems in the parameterized ensemble in (4) are available. Then, the sample moments satisfy $\mathbf{m}_k(t) \rightarrow m_k(t)$ for all $k \in \mathbb{N}$ as $q \rightarrow \infty$. Moreover, the convergence rate is bounded above by $O(q^{-1})$, provided that each Φ_k is (essentially) bounded.*

Proof Because the measured systems are randomly selected, β_i , $i = 1, \dots, q$ can be considered as a sequence of independent random variables, each of which has the uniform distribution on Ω . This indicates that $\Phi_k(\beta_i) x(t, \beta_i)$, $i = 1, \dots, q$, are independent and identically distributed random variables on \mathbb{R}^n . The expectation of each $\Phi_k(\beta_i) x(t, \beta_i)$, which is $\mathbb{E}(\Phi_k x_t) = \frac{1}{|\Omega|} \int \Phi_k(\beta) x_t(\beta) d\beta$, coincides with $\frac{1}{|\Omega|} m_k(t)$ and satisfies $\mathbb{E}|\Phi_k x_t| = \frac{1}{|\Omega|} \int \Phi_k(\beta) x_t(\beta) d\beta \leq \frac{1}{|\Omega|} \|\Phi_k\|_{\mathcal{H}} \|x_t\|_{\mathcal{H}} = \frac{1}{|\Omega|} \|x_t\|_{\mathcal{H}} < \infty$. The strong law of large numbers then implies $\mathbf{m}_k(t) \rightarrow m_k(t)$ almost surely as $q \rightarrow \infty$ (Billingsley, 1995).

To derive the bound of the convergence rate, it is necessary to compute the tail probability $\mathbb{P}(|\mathbf{m}_k(t) - m_k(t)| > \varepsilon)$. We claim that when Φ_k is essentially bounded by some constant C , then the random variable $\Phi_k x_t$ has a finite variance. This follows from $\mathbb{E}|\Phi_k x_t|^2 = \frac{1}{|\Omega|} \int \Phi_k^2(\beta) x_t^2(\beta) d\beta \leq \frac{C^2}{|\Omega|} \int x_t^2(\beta) d\beta = \frac{C^2 \|x_t\|_{\mathcal{H}}^2}{|\Omega|} < \infty$. By Chebyshev's inequality, we have

$$\begin{aligned} \mathbb{P}(|\mathbf{m}_k(t) - m_k(t)| > \varepsilon) &\leq \frac{\mathbb{E}|\mathbf{m}_k(t) - m_k(t)|^2}{\varepsilon^2} = \frac{|\Omega|^2}{(q\varepsilon)^2} \mathbb{E} \left| \sum_{i=1}^q \left(\Phi_k(\beta_i) x_t(\beta_i) - \frac{m_k(t)}{|\Omega|} \right) \right|^2 \\ &= \frac{|\Omega|^2}{(q\varepsilon)^2} \mathbb{E} \left| \sum_{i=1}^q \left(\Phi_k(\beta_i) x_t(\beta_i) - \mathbb{E}(\Phi_k x_t) \right) \right|^2 = \frac{|\Omega|^2}{(q\varepsilon)^2} \sum_{i=1}^q \mathbb{E} \left| \Phi_k(\beta_i) x_t(\beta_i) - \mathbb{E}(\Phi_k x_t) \right|^2 \\ &= \frac{|\Omega|^2}{(M\varepsilon)^2} \sum_{i=1}^q \text{Var}(\Phi_k x_t) \leq \frac{C^2 |\Omega| \|x_t\|_{\mathcal{H}}^2}{q\varepsilon^2} \sim O\left(\frac{1}{q}\right), \end{aligned}$$

where the second and third equalities follow from the independence of β_i , giving the desired bound. \blacksquare

Because the bound on the convergence rate is independent of the order of moments, the entire sample moment sequence $\mathbf{m}(t) = (\mathbf{m}_k(t))_{k \in \mathbb{N}}$ converges to the moment sequence

$m(t)$ with the rate $O(q^{-1})$ as well. In practice, we compute the sample moments up to a finite order N , constituting an order- N truncated sample moment sequence $\widehat{\mathbf{m}}_N(t)$. This computation only requires basic linear algebra operations with low computational complexity. To illustrate this, we concatenate the states of the q measured systems into a vector $\widehat{X}_q(t) = [x'(t, \beta_1), \dots, x'(t, \beta_q)]' \in \mathbb{R}^{qn}$ and define the *moment kernel matrix* $\widehat{\Psi}_N \in \mathbb{R}^{(N+1) \times M}$, whose (i, j) -entry is given by $\Psi_i(\beta_j)$. By the definition of sample moments in (19), the order- N truncated sample moment sequence, represented as a vector $\widehat{\mathbf{m}}_N(t) = [\mathbf{m}'_0(t), \dots, \mathbf{m}'_N(t)]' \in \mathbb{R}^{(N+1)n}$, is given by

$$\widehat{\mathbf{m}}_N(t) = (\widehat{\Psi}_N \otimes I_n) \cdot \widehat{X}_q(t), \quad (20)$$

where I_n denotes the n -by- n identity matrix and \otimes is the Kronecker product of matrices. In terms of computations, $\widehat{\Psi}_N \otimes I_n$ essentially arranges the $(N+1)q$ pre-determined numbers $\Psi_i(\beta_j)$ into an $(N+1)n$ -by- qn (block) matrix with $(N+1)q$ n -by- n diagonal blocks, which is of $O(1)$ complexity. The matrix multiplication in (20), in the worst case, requires $(N+1)qn$ scalar multiplications and $(N+1)(q-1)n$ additions, resulting in a complexity of $O(Nqn)$. In total, the complexity of computing the order- N truncated moment sequence using (20) is $O(NMn)$ in the worst-case scenario.

In addition, because the vector field \bar{F} of the moment kernelized system in (8) is the moment sequence of F from the parameterized ensemble system in (4), as discussed in Section 3.1, the above computational and sample efficiency analysis is directly applicable to learning the moment kernelized system.

4.1.2 SECOND-ORDER POLICY SEARCH

To showcase the FRL algorithm with early stopped hierarchies, in this section, we propose a second-order PS algorithm and integrate it into the FRL structure with an early stopping criterion. The development of the PS algorithm is based on the theory of differential dynamic programming (Jacobson and Mayne, 1970; Mayne, 1966). As a major advantage, the algorithm does not involve time discretization and is directly applicable to infinite-dimensional continuous-time systems, particularly the moment kernelized ensemble system in (8). The main idea is to expand the value function into a Taylor series up to the second-order term and then derive an update rule represented in the form of differential equations. In the following, we will highlight the key steps in the development of the algorithm (see Appendix C for the detailed derivation).

Quadratic approximation of Hamilton-Jacobi-Bellman equations. Let $\langle \cdot, \cdot \rangle : \mathcal{M} \times \mathcal{M} \rightarrow \mathbb{R}$ be the inner product on the RKHS \mathcal{M} . For any variation $\delta m \in \mathcal{M}$ at $m \in \mathcal{M}$, applying Taylor's theorem to the (\mathcal{M} -component of) value function $V^* : [0, T] \times \mathcal{M} \rightarrow \mathbb{R}$ yields $V^*(t, m + \delta m) = V^*(t, m) + \langle DV^*(t, m), \delta m \rangle + \frac{1}{2} \langle D^2 V^*(t, m) \cdot \delta m, \delta m \rangle + o(\delta m^2)$, where the ‘‘Hessian’’ $D^2 V^*(t, m)$ is identified with a bounded linear operator from $\mathcal{M} \rightarrow \mathcal{M}$, and $D^2 V^*(t, m) \cdot \delta m$ denotes the evaluation of $D^2 V^*(t, m)$ at δm (Lang, 1999). The quadratic approximation of the HJB equation in (14) is subsequently obtained by replacing V^* with

its Taylor expansion while neglecting the high-order term $o(\delta m^2)$, that is,

$$\begin{aligned} & \frac{\partial}{\partial t} V^*(t, m) + \left\langle \frac{\partial}{\partial t} DV^*(t, m), \delta m \right\rangle + \frac{1}{2} \left\langle \frac{\partial}{\partial t} D^2 V^*(t, m) \cdot \delta m, \delta m \right\rangle \\ & + \min_u \left\{ H(t, m + \delta m, u, DV^*(t, m)) + \left\langle D^2 V^*(t, m) \cdot \delta m, \bar{F}(t, m + \delta m, u) \right\rangle \right\} = 0, \end{aligned} \quad (21)$$

where $H(t, m + \delta m, u, DV^*(t, m)) = \bar{r}(t, m, u) + \langle DV^*(t, m), \bar{F}(t, m, u) \rangle$ is the *Hamiltonian* of the moment kernelized system.

Second-order policy search. The next step is to design an iterative algorithm to solve (21). The main idea is to generate a policy learning sequence $u^{(k)}$ using the update rule $u^{(k+1)} = u^{(k)} + \delta u^{(k)}$ such that $m^* = \lim_{k \rightarrow \infty} m^{(k)}$, where $m^{(k)}$ and m^* are the trajectories of the moment system steered by $u^{(k)}$ and u^* , respectively. Then, the (m^*, u^*) solves (21). To find $\delta u^{(k)}$, we note that because the pair $(m^{(k)}, u^{(k)})$ satisfies the moment system in (8), $\delta m^{(k)} = 0$ holds. Thus, the equation in (21) reduces to $\frac{\partial}{\partial t} V^*(t, m^{(k)}) + \min_u H(t, m^{(k)}, u, DV(t, m^{(k)})) = 0$, where the solution $\tilde{u}^{(k)} = \operatorname{argmin}_u H(t, m^{(k)}, u, DV(t, m^{(k)}))$ is expressed in terms of $m^{(k)}$ and $DV^*(t, m^{(k)})$. Steered by this policy $\tilde{u}^{(k)}$, the trajectory of the moment system is no longer $m^{(k)}$. Hence, a variation $\delta m^{(k)}$ on $m^{(k)}$ is produced such that $m^{(k)} + \delta m^{(k)}$ satisfies the moment system. This updates the equation in (21) to

$$\begin{aligned} & \frac{\partial}{\partial t} V^*(t, m^{(k)}) + \left\langle \frac{\partial}{\partial t} DV^*(t, m^{(k)}), \delta m^{(k)} \right\rangle + \frac{1}{2} \left\langle \frac{\partial}{\partial t} D^2 V^*(t, m^{(k)}) \cdot \delta m^{(k)}, \delta m^{(k)} \right\rangle \\ & + \min_{\delta u} \left\{ H(t, m^{(k)} + \delta m^{(k)}, \tilde{u}^{(k)} + \delta u, DV^*(t, m^{(k)})) \right. \\ & \left. + \left\langle D^2 V^*(t, m^{(k)}) \cdot \delta m^{(k)}, \bar{F}(t, m^{(k)} + \delta m^{(k)}, \tilde{u}^{(k)} + \delta u) \right\rangle \right\} = 0. \end{aligned} \quad (22)$$

To solve the minimization problem in (22), we expand the objective function into a Taylor series with respect to $(\delta m^{(k)}, \delta u)$ up to the second-order term, and then compute the critical point $\delta u^{(k)}$ in terms of $DV^*(t, m^{(k)})$ and $D^2 V^*(t, m^{(k)})$. Consequently, it remains to find $DV^*(t, m^{(k)})$ and $D^2 V^*(t, m^{(k)})$. Following the Taylor expansion, the equation in (22) becomes an algebraic second-order polynomial equation in $\delta m^{(k)}$, which is satisfied if all the coefficients are 0. This yields a systems of three ordinary differential equations, given by

$$\frac{d}{dt} \delta V(t, m^{(k)}(t)) = H(t, m^{(k)}, u^{(k)}, DV^*(t, m^{(k)})) - H(t, m^{(k)}, \tilde{u}^{(k)}, DV^*(t, m^{(k)})), \quad (23)$$

$$\frac{d}{dt} DV^*(t, m^{(k)}(t)) = -DH - D^2 V \cdot (\bar{F} - \bar{F}(t, m^{(k)}, u^{(k)})), \quad (24)$$

$$\begin{aligned} \frac{d}{dt} D^2 V^*(t, m^{(k)}(t)) &= -D^2 H - D\bar{F}' \cdot D^2 V^* - D^2 V^{*'} \cdot D\bar{F} \\ &+ \left[\frac{\partial DH}{\partial u} + \frac{\partial \bar{F}'}{\partial u} \cdot D^2 V^* \right]' \cdot \left(\frac{\partial^2 H}{\partial u^2} \right)^{-1} \cdot \left[\frac{\partial DH}{\partial u} + \frac{\partial \bar{F}'}{\partial u} \cdot D^2 V^* \right] \end{aligned} \quad (25)$$

with the terminal conditions $DV^*(T, m^{(k)}(T)) = D\bar{K}(T, m^{(k)}(T))$ and $D^2 V^*(T, m^{(k)}(T)) = D\bar{K}(T, m^{(k)}(T))$, where $\delta V(t, m^{(k)}(t)) = V^*(t, m^{(k)}(t)) - V(t, m^{(k)}(t))$ and $V(t, m^{(k)}(t))$ is the state-value function of the moment system driven by the policy $u^{(k)}$. Moreover, to simplify the notations, all the functions in the equations in (24) and (25) without arguments are evaluated at $(t, m^{(k)}, \tilde{u}^{(k)})$. The notation ' \prime ' denotes the transpose of linear operators.

Early stopping criterion. Because the development of the second-order PS algorithm is based on the Taylor series approximation, it is required that the amplitudes of $\delta m^{(k)}$ and $\delta u^{(k)}$ are small enough for all k . A necessary condition to guarantee this is to bound $\delta V^{(k)} = \sup_t |\delta V(t, m^{(k)}(t))|$ by a threshold η . When integrating the proposed second-order PS algorithm into each hierarchy of the FRL structure, say the i^{th} hierarchy, the algorithm will be terminated if $\sup_t |\hat{V}_{N_i}(t, \hat{m}_{N_i}^{(k)}(t)) - \hat{V}_{N_i}(t, \hat{m}_{N_i}^{(k-1)}(t))| > \eta$. The resulting policy $u^{(k)}$ will then be used as the initial condition for the $(i+1)^{\text{th}}$ hierarchy. The FRL algorithm with the early stopped second-order policy search is shown in Algorithm 2.

Algorithm 2 Filtrated reinforcement learning for parameterized systems with early stopped second-order policy search hierarchies

Input: Initial state x_0 , final time T , projection error tolerance ε , value function variation tolerance η , maximum number of policy search iterations K

2: *Initialization:* truncation order N_0 , control policy $u^{(0)}$, projection error $P > \varepsilon$, value function variation $\delta V = 0$, hierarchy level $i = 0$, policy search iteration number $j = 0$

while $P > \varepsilon$ **do**

4: **while** $\delta V \leq \eta$ **and** $j \leq K$ **do**

 Generate data by solving the parameterized system with the input $u^{(j)}$ and collect the rewards

6: Compute order N_i truncated moment kernelization: $\hat{m}_{N_i}(t)$, $\hat{F}_{N_i}(t, \hat{m}_{N_i}(t), u^{(0)})$, $\bar{r}(t, \hat{m}_{N_i}(t), u^{(0)})$, $\bar{K}(T, \hat{m}_{N_i}(T))$

 Compute $\delta \hat{V}_{N_i}(t, \hat{m}_{N_i}^{(j)}(t))$, $D\hat{V}_{N_i}^*(t, \hat{m}_{N_i}^{(j)}(t))$ and $D^2\hat{V}_{N_i}^*(t, \hat{m}_{N_i}^{(j)}(t))$ for $0 \leq t \leq T$ by solving the system of differential equations in (23), (24) and (25)

8: $u^{(j+1)}(t) \leftarrow \operatorname{argmin}_a H(t, \hat{m}_{N_i}^{(j)}(t), a, D\hat{V}_{N_i}(t, \hat{m}_{N_i}^{(j)}(t)))$ for all $0 \leq t \leq T$

$\delta V \leftarrow \max_t |\delta \hat{V}(t, \hat{m}_{N_i}^{(j)})|$

10: $j \leftarrow j + 1$

end while

12: $\hat{m}_{N_i}^* \leftarrow \hat{m}_{N_i}^{(j-1)}$

if $i \geq 1$ **then**

14: $P = \max_{t \in [0, T]} |\hat{V}_{N_i}(t, \hat{m}_{N_i}^*(t)) - \hat{V}_{N_{i-1}}(t, \hat{m}_{N_{i-1}}^*(t))|$

end if

16: $N_{i+1} \leftarrow N_i + 1$, $u^{(0)} \leftarrow u^{(j-1)}$

$i \leftarrow i + 1$

18: **end while**

$u^* \leftarrow u^{(j-1)}$ **return** u^*

4.2 Examples and simulations

In this section, we will demonstrate the performance and efficiency of the proposed FRL algorithm using parameterized ensemble systems arising from practical applications. All the simulations were run on an Apple M1 chip with 16 GB memory.

4.2.1 INFINITE-DIMENSIONAL LINEAR-QUADRATIC REGULATORS

Linear-quadratic (LQ) problems, those are, linear systems with state-value functions being quadratic in both the system states and control variables, are among the most fundamental control problems, which have been extensively studied for finite-dimensional linear systems (Brockett, 2015; Kwakernaak and Sivan, 1972; Sontag, 1998). However, LQ problems for parameterized ensemble systems defined on infinite-dimensional function spaces remain largely unexplored. In this example, we will fill in this literature gap to approach such type of infinite-dimensional LQ problems using the proposed FRL algorithm.

To illuminate the main idea as well as to demonstrate how FRL addresses the curse of dimensionality, we revisit the parameterized scalar linear system in (3), i.e. $\frac{d}{dt}x(t, \beta) = \beta x(t, \beta) + u(t)$, $\beta \in [-1, 1]$, with a finite-time horizon state-value function given by $V(t, x_t) = \int_{-1}^1 \left[\int_t^T (x^2(t, \beta) + u^2(t)) dt d\beta + x^2(T, \beta) \right] d\beta$. To apply the moment kernel transform, we consider the case where the system evolves on the Hilbert space $L^2([0, 1], \mathbb{R})$ and choose the basis $\{\Phi_k\}_{k \in \mathbb{N}}$ to be the set of Chebyshev polynomials. In this case, the moment kernelized system and state-value function can be analytically derived as $\frac{d}{dt}m(t) = Am(t) + bu(t)$ and $V(t, m(t)) = \int_0^T [\|m(t)\|^2 + 2u^2(t)] dt + \|m(T)\|^2$, where $A = \frac{1}{2}(L + R)$ with L and R the left- and right-shift operators, given by $(m_0(t), m_1(t), \dots) \mapsto (m_1(t), m_2(t), \dots)$ and $(m_0(t), m_1(t), \dots) \mapsto (0, m_0(t), \dots)$, respectively (see Appendix D.1 for the detailed derivation). However, the analytic form of the moment kernelized system and state-value function are not required in the implementation of FRL.

In the simulation, the final time and initial condition for the parameterized system were set to $T = 1$ and $x_0 = 1$, the constant function on $[-1, 1]$, respectively, and the tolerance for the value function variation at each hierarchy was chosen to be $\eta = 1$. The truncation order N was varied from 2 to 10, and in each case the evolution of the truncated moment kernelized system was approximated using the sample moments computed from the measurement data for 500 systems with their system parameters β uniformly sampled from $[-1, 1]$. The simulation results generated by Algorithm 2 are shown in Figure 3. Specifically, Figure 3a shows the total reward $\hat{V}(0, \hat{m}_N(0))$ (top panel) and the number of the policy search iterations (bottom panel) with respect to the truncation order (hierarchy level) N . We observe that the total reward converges to the minimum value after only 4 hierarchies of the algorithm, which demonstrates the high efficiency of FRL. Correspondingly, Figure 3b plots the policy u_N learned from each hierarchy of FRL, which stabilizes in the shadowed region starting from $N = 6$. It is also worth mentioning that the computational time for running 10 hierarchies of the algorithm is only 3.97 seconds, indicating the low computational cost of FRL. As a result, the curse of dimensionality is effectively mitigated.

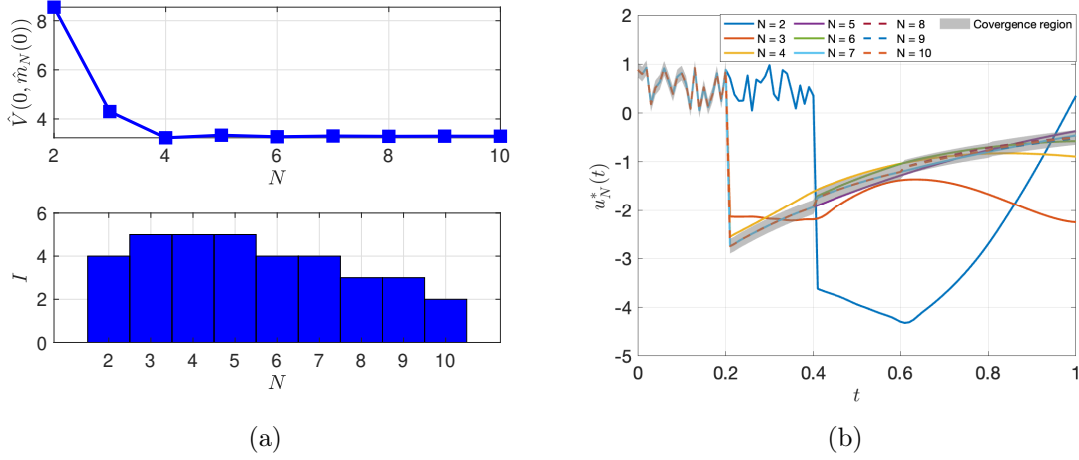


Figure 3: Learning the finite-time horizon LQR policy for the parameterized ensemble system in (3) using FRL with early stopped hierarchies shown in Algorithm 2. In particular, (a) shows the total reward (top panel) and the number of the second-order PS iterations with respect to the truncation order (hierarchy level) N ranging from 2 to 20, and (b) plots the learned policy from each hierarchy of FRL.

To further demonstrate the advantages of FRL, we will show that it also resolves the convergence issue caused by applying classical RL algorithms to sampled parameterized systems, as discussed in Section 2.1. To this end, we revisit the infinite-time horizon LQR problem presented there, that is, the same system in (3) with the state-value function given by $V(x_t) = \int_{-1}^1 \int_0^\infty e^{-2.5t} [x^2(t, \beta) + u^2(t)] dt d\beta$. In this case, we use Algorithm 1 with the standard value iteration applied to each hierarchy, and the simulation results are shown in Figure 4. A comparison between Figures 4a and 1b reveals that now the projection error for both the learned value functions and optimal policies converge to 0. This means that the sequences of value functions and optimal policies generated by FRL in Algorithm 1 are Cauchy sequences, and hence necessarily converge to those of the parameterized system in (3) (Rudin, 1976). As further verification, Figure 4b illustrates that the learned value functions and optimal policies stabilize in the corresponding shadowed regions.

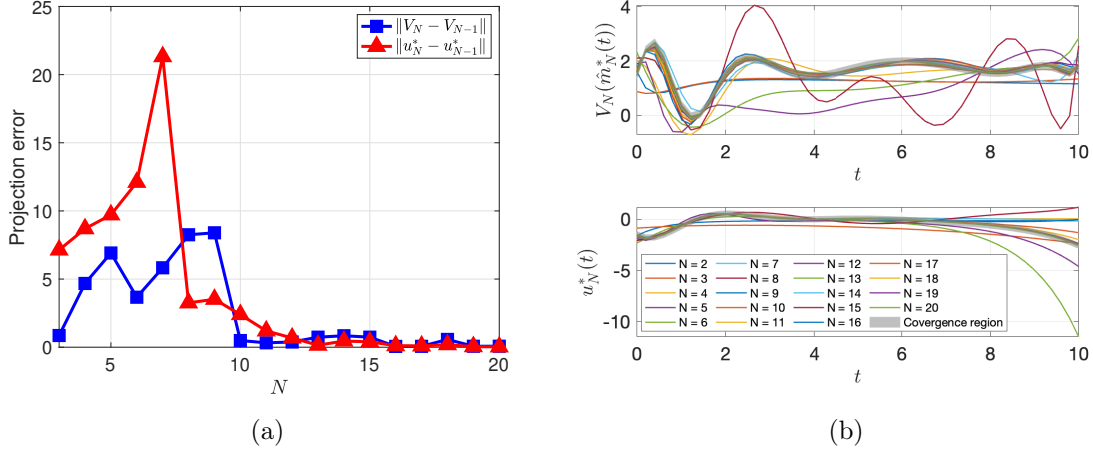


Figure 4: FRL resolves the convergence issue caused by applying classical RL algorithms to sampled parameterized systems. Algorithm 1 with the standard value iteration applied to each hierarchy is used to learn the infinite-time horizon LQR policy and value function of the parameterized system in (3). Specifically, (a) plots the projection error of the learned value functions and optimal policies, and (b) shows the learned value function (top) optimal policy (bottom) at each hierarchy.

4.2.2 FILTRATED REINFORCEMENT LEARNING FOR PARAMETERIZED NONLINEAR ENSEMBLE SYSTEMS

In this section, we apply FRL to learn the optimal policy for a parameterized nonlinear system arising from quantum mechanics and quantum control. The policy learning problem here is the robust excitation of a nuclear spin sample, typically consisting of as many spins as the order of Avogadro’s number ($\sim 10^{23}$). This problem, also known as the pulse design problem, is crucial in quantum science and technology. For example, it enables all the applications of nuclear magnetic resonance (NMR) spectroscopy, including magnetic resonance imaging (MRI), quantum computing, quantum optics, and quantum information processing (Li and Khaneja, 2009; Li et al., 2022; Silver et al., 1985; Roos and Moelmer, 2004; Stefanatos and Li, 2011; Chen et al., 2011; Stefanatos and Li, 2014). The dynamics of nuclear spins immersed in a static magnetic field, with respect to the rotating frame, is governed by the Bloch equation

$$\frac{d}{dt}x(t, \beta) = \begin{bmatrix} 0 & 0 & -\beta u(t) \\ 0 & 0 & \beta v(t) \\ \beta u(t) & -\beta v(t) & 0 \end{bmatrix} x(t, \beta), \quad (26)$$

which was derived by the Swiss-American physicist Felix Bloch in 1946 (Cavanagh et al., 2010). In this system, the state variable $x(t, \beta)$ denotes the bulk magnetization of spins, $u(t)$ and $v(t)$ represent the external radio frequency (rf) fields, and the system parameter $\beta \in \Omega = [1 - \delta, 1 + \delta]$ with $0 < \delta < 1$ is referred to as the *rf inhomogeneity*, characterizing the phenomenon that spins in different positions of the sample receive different strength of the rf fields. In practice, the inhomogeneity can be up to $\delta = 40\%$ of the strength of the applied rf fields (Nishimura et al., 2001). A typical policy learning task is to design the rf fields $u(t)$ and $v(t)$, with minimum energy, to steer the parameterized Bloch system in (26) from the *equilibrium state* $x_0(\beta) = (0, 0, 1)'$ to the *excited state* $x_F(\beta) = (1, 0, 0)'$ for all β . We formulate

this task as an RL problem over the function space $L^2(\Omega, \mathbb{R}^3)$, i.e., $x(t, \cdot) \in L^2(\Omega, \mathbb{R}^3)$, with the state-value function defined by $V(t, x_t) = \int_t^T (u^2(t) + v^2(t)) dt + \int_{1-\delta}^{1+\delta} |x(T, \beta) - x_F(\beta)|^2 d\beta$, where $|\cdot|$ denotes a norm on \mathbb{R}^3 .

We still use the Chebyshev polynomial basis $\{\Phi_k\}_{k \in \mathbb{N}}$ to define the moment transform for kernelizing the Bloch system and the state-value function. To guarantee the orthonormality of $\{\Phi_k\}_{k \in \mathbb{N}}$, we first rescale the range of β from $[1-\delta, 1+\delta]$ to $[-1, 1]$ by the linear transformation $\beta \mapsto \eta = (\beta - 1)/\delta$, and then apply the moment kernel transform defined in (6) (see Appendix D.2 for the detailed derivation). We then apply Algorithm 2 to learn the optimal policies for the spin system in (26). In the simulation, we considered the maximal rf inhomogeneity $\delta = 40\%$ encountered in practice and pick the final time to be $T = 1$. Similar to the previous example, we varied the moment truncation order N from 2 to 10; for each N we approximated the evolution of the truncated moment kernelized system using the sample moments computed from the measurement data for 500 systems in the ensemble with the system parameters uniformly sampled from $[0.6, 1.4]$. The results are shown in Figure 5. In particular, Figure 5a displays the total reward (top panel) and number of second-order policy search iterations (bottom panel) with respect to the moment truncation order N , that is, the hierarchy level, from which we observe their convergence at $N = 9$. Figure 5b plots the policy and value function learned from each hierarchy, which stabilize in the corresponding shadowed regions as the truncation order increases. The computational time is 42.30 seconds, which is much longer than that (3.97 seconds) for the LQR problem presented in Section 4.2.1. In part, this is because the order- N truncated moment system in this case is of dimension $3N$, 3 times higher than the moment system in the LQR example. Additionally, the nonlinearity of the Bloch system increases the complexity of this policy learning task.

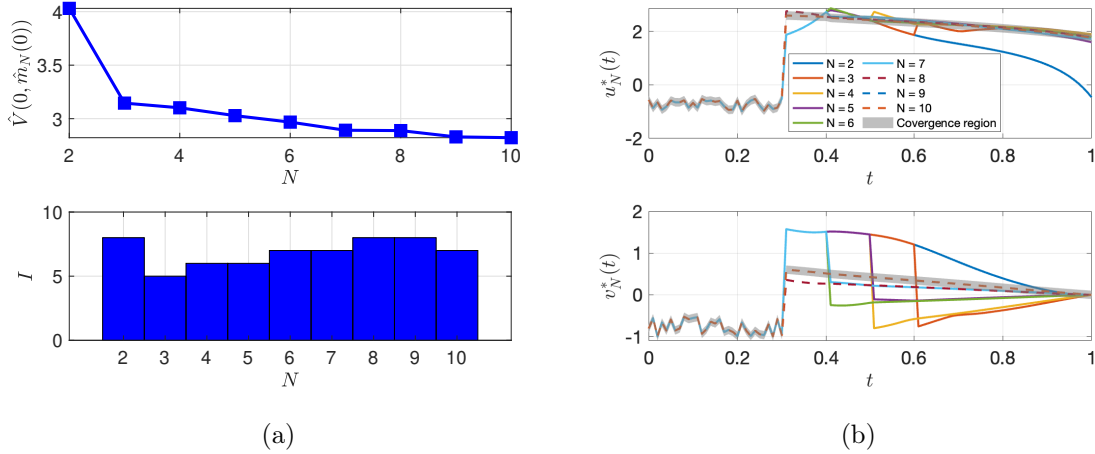


Figure 5: Policy learning for robust excitation of the nuclear spin system in (26) in the presence of rf inhomogeneity using FRL shown in Algorithm 2. In particular, (a) shows the total reward (top panel) and number of the second-order PS iterations with respect to the moment truncation order, that is, the hierarchy level, N ranging from 2 to 10, and (b) plots the policies learned from each hierarchy of the algorithm.

As mentioned previously, from the perspective of quantum physics, a goal of controlling the parameterized Bloch system is to steer the spins from the equilibrium state $(0, 0, 1)'$ to the excited state $(1, 0, 0)'$ uniformly regardless of the rf inhomogeneity. Note that because

$\Omega_x = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & -1 & 0 \end{bmatrix}$ and $\Omega_y = \begin{bmatrix} 0 & 0 & -1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}$ in the Bloch system $\frac{d}{dt}x(t, \beta) = \beta[u(t)\Omega_y + v(t)\Omega_x]x(t, \beta)$ are skew-symmetric matrices, we have $\frac{d}{dt}|x(t, \beta)|^2 = \frac{d}{dt}(x'(t, \beta)x(t, \beta)) = \frac{d}{dt}x'(t, \beta) \cdot x(t, \beta) + x'(t, \beta) \cdot \frac{d}{dt}x(t, \beta) = x'(t, \beta)(\beta u\Omega'_y + v\Omega'_x)x(t, \beta) + x'(t, \beta)(\beta u\Omega_y + v\Omega_x)x(t, \beta) = 0$. As a result, starting from $x(0, \beta) = (0, 0, 1)$, $|x(t, \beta)| = 1$ holds for all t and β , meaning the system trajectory stays on the unit sphere $\mathbb{S}^2 = \{x \in \mathbb{R}^3 : |x| = 1\}$. To evaluate the excitation performance, it then suffices to examine $x_1(T, \beta)$ solely, the first component of the final state $x(T, \beta)$, which is plotted in the top panel of Figure 6a as a function of β . Its average value, generally used as the measure of the control performance (Zhang and Li, 2015), is $\frac{1}{2\delta} \int_{1-\delta}^{1+\delta} x_1(T, \beta) d\beta = 0.9613$, which is close to 1, indicating the good excitation performance of the learned policies. The bottom panel of Figure 6a shows the performance measure versus time. Figure 6b shows the entire trajectory of the Bloch system on \mathbb{S}^2 , from which we observe that the learned policies indeed steer the system towards the excited state $(1, 0, 0)'$, regardless of the rf inhomogeneity, as desired.

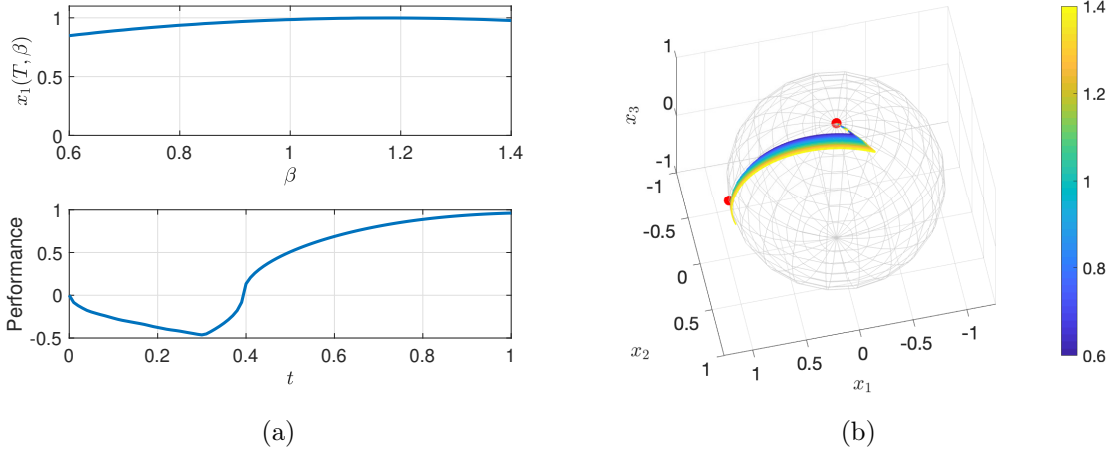


Figure 6: Excitation performance of the policies learned by using Algorithm 2. In particular, (a) plots the first component of the final state of the parameterized Bloch system in (26) with respect to the rf inhomogeneity β , and (b) plots the entire trajectory of the system driven by the learned policies on the unit sphere.

4.2.3 COMPARISON WITH BASELINE DEEP REINFORCEMENT LEARNING MODELS

We conduct a comparative analysis between the proposed FRL architecture with baseline deep RL models. Given our problem formulation, where policies are deterministic and take values in continuous action spaces, we choose the Deep Deterministic Policy Gradient (DDPG) and Twin-Delayed Deep Deterministic Policy Gradient (TD3) as the baseline models. We test their performance on 500 systems in the parametrized linear ensemble in (3) as well as 500 systems in the parametrized Bloch ensemble in (26), with their system parameters randomly sampled from $[-1, 1]$ and $[0.6, 1.4]$, respectively. In table 1, we compare the minimal costs and training times of DDPG and TD3 with those of the proposed FRL. We observe that FRL outperforms the baseline deep RL models in both the linear and Bloch system

examples. Notably, FRL achieves significant computational efficiency, reducing the training time compared to deep RL models by two orders of magnitude.

Linear system			
	FRL	DDPG	TD3
Minimal cost	3.40	59.86	50.35
Training time (s)	3.97	1607	1575
Bloch (nonlinear) system			
	FRL	DDPG	TD3
Minimal cost	2.85	28.95	32.46
Training time (s)	42.30	2807	2240

Table 1: Comparison of the proposed FRL with baseline deep RL models, DDPG and TD3. The top and bottom panels show the comparison results for the parameterized linear system in (3) and (nonlinear) Bloch system in (26), respectively.

In addition, a major drawback of the deep RL models lies in their failure to retain the geometric structure of the Bloch system. Recall that the state of the parameterized Bloch system should stay on the unit sphere \mathbb{S}^2 . However, as shown in Figure 7, neither the trajectories learned from DDPG and TD3 satisfy this requirement. This indicates that both the DDPG and TD3 agents fail to learn the evolution of the Bloch system.

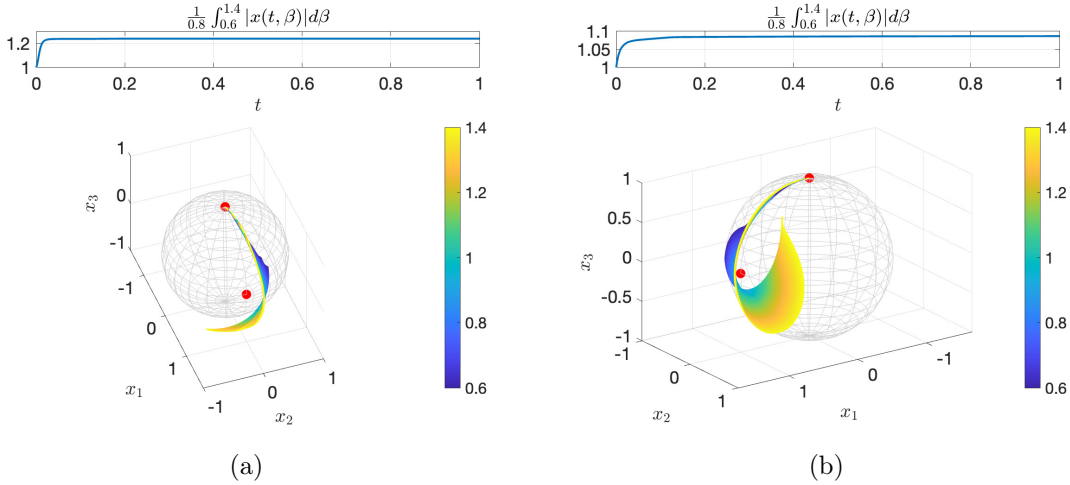


Figure 7: Temporal evolution of the parameterized Bloch system in (26) learned by using DDPG (a) and TD3 (b). In both (a) and (b), the bottom panels show the the learned trajectories and the top panels plot the “averaged” norms of the trajectories over the entire ensemble.

5. Conclusion

In this paper, we propose a novel RL architecture designed to learn optimal policies for arbitrarily large populations of intelligent agents. Our approach models such populations as parameterized control systems defined on infinite-dimensional function spaces. To address the inherent challenges posed by infinite-dimensionality, we develop the moment kernel transform,

which maps the parameterized system and its value function to an RKHS consisting of moment sequences. This transformation results in a kernel parameterization of the RL problem. We then structure the finite-dimensional truncated moment representations of the RL problem into a filtrated framework and develop a hierarchical policy learning algorithm through this RL filtration. Each hierarchy in this framework corresponds to an RL problem for a finite-dimensional truncated moment kernelized system. To enhance the computational efficiency of the FRL algorithm, we investigate early stopping criteria for each hierarchy and establish the convergence of the early-stopped algorithm by constructing a spectral sequence. Additionally, we quantitatively analyze the computational and sample efficiency of the FRL. Examples are provided to demonstrate the exceptional performance and high efficiency of the proposed algorithm, validating its effectiveness in practical applications.

Acknowledgments

We are grateful to the anonymous reviewers and the AE for their insightful comments and suggestions, which have helped improve the presentation of this paper. This work was supported by the Air Force Office of Scientific Research under award FA9550-21-1-0335.

Appendix A. Proof of Lemma 2

Let $x_k(t, \cdot)$ and $x(t, \cdot)$ in $\mathcal{F}(\Omega, M)$ be the trajectories (solutions) of the ensemble system in (4), with the initial condition $x(0, \cdot)$, driven by the control inputs u_k and u , respectively, for all $k \in \mathbb{N}$, i.e.,

$$\frac{d}{dt}x_k(t, \beta) = F(t, \beta, x_k(t, \beta), u_k(t)) \quad \text{and} \quad \frac{d}{dt}x(t, \beta) = F(t, \beta, x(t, \beta), u(t)).$$

We claim that $x_k(t, \cdot)$ converges to $x(t, \cdot)$ on $\mathcal{F}(\Omega, M)$ pointwisely as $k \rightarrow \infty$ for any $t \in [0, T]$, which is equivalent to $x_k(t, \beta) \rightarrow x(t, \beta)$ for any $\beta \in \Omega$ on each coordinate chart of M by using a partition of unity on M (Lee, 2012). Therefore, it suffices to assume that the entire trajectories $x_k([0, T], \beta) \subseteq M$ and $x([0, T], \beta) \subseteq M$ for all $k \in \mathbb{N}$ and each $\beta \in \Omega$ are located in a single coordinate chart of M and hence equivalently in \mathbb{R}^n .

To prove the claim, we fix an arbitrary $\beta \in \Omega$ and note that

$$\begin{aligned} \frac{d}{dt}(x_k(t, \beta) - x(t, \beta)) &= F(t, \beta, x_k(t, \beta), u_k(t)) - F(t, \beta, x(t, \beta), u(t)) \\ &\leq \left| F(t, \beta, x_k(t, \beta), u_k(t)) - F(t, \beta, x(t, \beta), u(t)) \right| \\ &\leq \left| F(t, \beta, x_k(t, \beta), u_k(t)) - F(t, \beta, x(t, \beta), u_k(t)) \right| \\ &\quad + \left| F(t, \beta, x(t, \beta), u_k(t)) - F(t, \beta, x(t, \beta), u(t)) \right|, \end{aligned}$$

where the second inequality follows from the triangle inequality. Because $\frac{d}{dt}(x(t, \beta) - x_k(t, \beta))$ satisfies the same inequality as above, we obtain

$$\begin{aligned} \frac{d}{dt}|x_k(t, \beta) - x(t, \beta)| &\leq \left| F(t, \beta, x_k(t, \beta), u_k(t)) - F(t, \beta, x(t, \beta), u_k(t)) \right| \\ &\quad + \left| F(t, \beta, x(t, \beta), u_k(t)) - F(t, \beta, x(t, \beta), u(t)) \right| \\ &\leq C|x_k(t, \beta) - x(t, \beta)| + \left| F(t, \beta, x(t, \beta), u_k(t)) - F(t, \beta, x(t, \beta), u(t)) \right| \end{aligned}$$

where the second inequality follows from the Lipschitz continuity of F in the system state variable according to Assumption S2. Moreover, as solutions of ordinary differential equations, all of $x_k(t, \beta)$ and $x(t, \beta)$ are Lipschitz continuous, and hence absolutely continuous, then so is $|x_k(t, \beta) - x(t, \beta)|$. Together with its nonnegativity, Gronwall's inequality can be applied (Evans, 2010), yielding

$$|x_k(t, \beta) - x(t, \beta)| \leq e^{Ct} \int_0^t \left| F(s, \beta, x(s, \beta), u_k(s)) - F(s, \beta, x(s, \beta), u(s)) \right| ds$$

Because F is continuous in the control policy variable, the pointwise convergence of $u_k(s)$ to $u(s)$ implies that of $F(s, \beta, x(s, \beta), u_k(s))$ to $F(s, \beta, x(s, \beta), u(s))$. Then, by Egoroff's Theorem, there exist sequences of real numbers $\varepsilon_k > 0$ and subsets I_k of $[0, T]$ with Lebesgue measure ε_k such that $\varepsilon_k \rightarrow 0$ and $F(s, \beta, x(s, \beta), u_k(s)) \rightarrow F(s, \beta, x(s, \beta), u(s))$ uniformly

on $I_k^c = [0, T] \setminus I_k$. Consequently, we have

$$\begin{aligned}
 \lim_{k \rightarrow \infty} |x_k(t, \beta) - x(t, \beta)| &\leq \lim_{k \rightarrow \infty} e^{Ct} \int_0^t \left| F(s, \beta, x(s, \beta), u_k(s)) - F(s, \beta, x(s, \beta), u(s)) \right| ds \\
 &\leq e^{Ct} \left\{ \lim_{k \rightarrow \infty} \int_{I_k^c} \left| F(s, \beta, x(s, \beta), u_k(s)) - F(s, \beta, x(s, \beta), u(s)) \right| ds \right. \\
 &\quad \left. + \lim_{k \rightarrow \infty} \int_{I_k} \left| F(s, \beta, x(s, \beta), u_k(s)) - F(s, \beta, x(s, \beta), u(s)) \right| ds \right\} \\
 &= e^{Ct} \left\{ \int_0^t \lim_{k \rightarrow \infty} \left| F(s, \beta, x(s, \beta), u_k(s)) - F(s, \beta, x(s, \beta), u(s)) \right| \chi_{I_k^c}(s) ds \right. \\
 &\quad \left. + \lim_{k \rightarrow \infty} \int_{I_k} \left| F(s, \beta, x(s, \beta), u_k(s)) - F(s, \beta, x(s, \beta), u(s)) \right| ds \right\} = 0,
 \end{aligned}$$

in which $\chi_{I_k^c}$ denotes the characteristic function of I_k^c , i.e., $\chi_{I_k^c}(s) = 1$ for $s \in I_k^c$ and $\chi_{I_k^c}(s) = 0$ for $s \notin I_k^c$, the change of the limit and integral in the first term in the summand follows from the uniform convergence of $F(s, \beta, x(s, \beta), u_k(s))$ to $F(s, \beta, x(s, \beta), u(s))$ on I_k^c , and the second integral converges 0 because the Lebesgue of I_k goes to 0. This then proves the claim.

Now, without loss of generality, we assume that $|J(u)| < \infty$, and the existence of such a u is guaranteed by Assumption C1. Then, we obtain the desired convergence

$$\begin{aligned}
 \lim_{k \rightarrow \infty} J(u_k) &= \lim_{k \rightarrow \infty} \left[\int_{\Omega} \int_0^T r(t, x_k(t, \beta), u_k(t)) dt d\beta + \int_{\Omega} K(T, x_k(T, \beta)) d\beta \right] \\
 &= \int_{\Omega} \int_0^T \lim_{k \rightarrow \infty} r(x_k(t, \beta), u_k(t)) dt d\beta + \int_{\Omega} \lim_{k \rightarrow \infty} K(x_k(T, \beta)) d\beta \\
 &= \int_{\Omega} \int_0^T r(t, \lim_{k \rightarrow \infty} x_k(t, \beta), \lim_{k \rightarrow \infty} u_k(t)) dt d\beta + \int_{\Omega} K(T, \lim_{k \rightarrow \infty} x_k(T, \beta)) d\beta \\
 &= \int_{\Omega} \int_0^T r(t, x(t, \beta), u(t)) dt d\beta + \int_{\Omega} K(T, x(T, \beta)) d\beta = J(u),
 \end{aligned}$$

where the second and third equalities follow from the dominated convergence theorem and the continuity of r and K , respectively (Folland, 2013).

Appendix B. Proof of Proposition 8

Fix $z, \tilde{z} \in \mathcal{M}$ and $t, \tilde{t} \in [0, T]$, then by the definition of the infimum, for any $\varepsilon > 0$, there exists an ensemble control policy $u \in \mathcal{U}$ such that

$$\begin{aligned}
 V^*(\tilde{t}, \tilde{z}) + \varepsilon &\geq \int_{\tilde{t}}^T \bar{r}(\tau, \tilde{m}(\tau), u(\tau)) d\tau + \bar{K}(T, \tilde{m}(T)) \\
 &= \int_{\Omega} \int_{\tilde{t}}^T r(\tau, \tilde{x}(\tau, \beta), u(\tau)) d\tau d\beta + \int_{\Omega} K(T, \tilde{x}(T, \beta)) d\beta,
 \end{aligned}$$

where $\tilde{m}(\tau)$ satisfies the moment system $\frac{d}{d\tau} \tilde{m}(\tau) = \bar{F}(\tau, \tilde{m}(\tau), u(\tau))$ with the initial condition $\tilde{m}(s) = \tilde{z}$ and $\tilde{x}(t, \cdot)$ is the trajectory of the associated ensemble system. Next, let $m(\tau)$ be the

trajectory of the moment system steered by the same control policy with the initial condition $m(t) = z$ and $x(t, \cdot)$ be the associated ensemble trajectory. Without loss of generality, we assume that $t \leq \tilde{t}$, and then we have

$$\begin{aligned}
 V^*(t, z) - V^*(\tilde{t}, \tilde{z}) &\leq \int_{\Omega} \int_t^T r(\tau, x(\tau, \beta), u(\tau)) d\tau d\beta + \int_{\Omega} K(T, x(T, \beta)) d\beta \\
 &\quad - \int_{\Omega} \int_{\tilde{t}}^T r(\tau, \tilde{x}(\tau, \beta), u(\tau)) d\tau d\beta - \int_{\Omega} K(T, \tilde{x}(T, \beta)) d\beta + \varepsilon \\
 &= \int_{\Omega} \int_t^{\tilde{t}} r(\tau, x(\tau, \beta), u(\tau)) d\tau d\beta + \int_{\Omega} \int_{\tilde{t}}^T [r(\tau, x(\tau, \beta), u(\tau)) \\
 &\quad - r(\tau, \tilde{x}(\tau, \beta), u(\tau))] d\tau d\beta + \int_{\Omega} [K(T, x(T, \beta)) - K(T, \tilde{x}(T, \beta))] d\beta + \varepsilon.
 \end{aligned} \tag{27}$$

Because finite-time solutions of ordinary differential equations with initial conditions in a bounded set remain bounded, $|r(\tau, x(\tau, \beta), u(\tau))| \leq M$ holds for some M over $\tau \in [t, \tilde{t}]$ so that the first term in (27) can be bounded as

$$\int_{\Omega} \int_t^{\tilde{t}} r(\tau, x(\tau, \beta), u(\tau)) d\tau d\beta \leq \int_{\Omega} \int_t^{\tilde{t}} |r(\tau, x(\tau, \beta), u(\tau))| d\tau d\beta M \leq \text{Vol}(\Omega) |t - \tilde{t}|.$$

Then, by using the Lipschitz continuity of r and K as in Assumption C2, the second and third terms in (27) satisfies

$$\begin{aligned}
 &\int_{\Omega} \int_{\tilde{t}}^T [r(\tau, x(\tau, \beta), u(\tau)) - r(\tau, \tilde{x}(\tau, \beta), u(\tau))] d\tau d\beta + \int_{\Omega} [K(T, x(T, \beta)) - K(T, \tilde{x}(T, \beta))] d\beta \\
 &\leq \int_{\Omega} \int_{\tilde{t}}^T C_r |x(\tau, \beta) - \tilde{x}(\tau, \beta)| d\tau d\beta + \int_{\Omega} C_K |x(T, \beta) - \tilde{x}(T, \beta)| d\beta \\
 &\leq \int_{\Omega} (TC_r + C_K) C' |x(\tilde{t}, \beta) - \tilde{x}(\tilde{t}, \beta)| d\beta \\
 &\leq (TC_r + C_K) C' \int_{\Omega} [|x(t, \beta) - \tilde{x}(\tilde{t}, \beta)| + |x(t, \beta) - x(\tilde{t}, \beta)|] d\beta \\
 &\leq (TC_r + C_K) C' \left[\int_{\Omega} |x(t, \beta) - \tilde{x}(\tilde{t}, \beta)| d\beta + \int_{\Omega} |x(t, \beta) - x(\tilde{t}, \beta)| d\beta \right],
 \end{aligned} \tag{28}$$

where the second inequality follows from Gronwall's inequality. Note that the first integral in (28) is exactly the L^1 -norm of $x(t, \cdot) - \tilde{x}(t, \cdot)$, which is equal to $\|z - \tilde{z}\|$, the norm of the associated moment sequences, since the moment transformation is an isometry. For the second term, we use the Lipschitz continuity of solutions to ordinary differential equations to conclude $\int_{\Omega} |x(t, \beta) - x(\tilde{t}, \beta)| \leq C'' \text{Vol}(\Omega) |t - \tilde{t}|$ for some $C'' > 0$. Now, let $C = \max\{(TC_r + C_K) C', M \text{Vol}(\Omega), \text{Vol}(\Omega)\}$, then we obtain $V^*(t, z) - V^*(\tilde{t}, \tilde{z}) \leq C(|t - \tilde{t}| + \|z - \tilde{z}\|)$ since $\varepsilon > 0$ is arbitrary. The same argument with the roles of (t, z) and (\tilde{t}, \tilde{z}) reversed implies

$$|V^*(t, z) - V^*(\tilde{t}, \tilde{z})| \leq C(|t - \tilde{t}| + \|z - \tilde{z}\|),$$

giving the Lipschitz continuity of V^* on $[0, T] \times \mathcal{M}$ as desired.

Appendix C. Derivation of Second-Order Policy Search Update Equations

Quadratic approximation of Hamilton-Jacobi-Bellman equations. We pick a nominal policy $\bar{u}(t)$, which generates a nominal trajectory $\bar{m}(t)$ by applying $\bar{u}(t)$ to the moment system in (8). Then, the optimal policy and trajectory can be represented as $u^*(t) = \bar{u}(t) + \delta u(t)$ and $m^*(t) = \bar{m}(t) + \delta m(t)$, respectively, plugging which into the Hamiltonian-Jacobi-Bellman equation in (14), i.e.,

$$\frac{\partial}{\partial t} V(t, z) + \min_{a \in U} \{ \bar{r}(t, z, a) + \langle DV(t, z), \bar{F}(t, z, a) \rangle \} = 0$$

yields

$$\begin{aligned} \frac{\partial}{\partial t} V(t, \bar{m} + \delta m) + \min_{\delta u} \{ \bar{r}(t, \bar{m} + \delta m, \bar{u} + \delta u) \\ + \langle DV(t, \bar{m} + \delta m), \bar{F}(t, \bar{m} + \delta m, \bar{u} + \delta u) \rangle \} = 0, \end{aligned}$$

where, and in the following as well, we drop the time argument t from the state and control variables for the conciseness of the representation. Now, we assume that the value function $V : [0, T] \times \mathcal{M} \rightarrow \mathbb{R}$ is smooth enough, at least in the region containing the nominal and optimal trajectories, to admit a second-order power series expansion as

$$\begin{aligned} V(t, \bar{m} + \delta m) &= V(t, \bar{m}) + \langle DV(t, \bar{m}), \delta m \rangle + \frac{1}{2} \langle D^2 V(t, \bar{m}) \cdot \delta m, \delta m \rangle + o(\delta m^2) \\ &= \bar{V}(t, \bar{m}) + \delta V(t, \bar{m}) + \langle DV(t, \bar{m}), \delta m \rangle + \frac{1}{2} \langle D^2 V(t, \bar{m}) \cdot \delta m, \delta m \rangle + o(\delta m^2), \end{aligned} \quad (29)$$

where we further expand $V(t, \bar{m})$ by using the nominal cost $\bar{V}(t, \bar{m}) = \int_t^T \bar{r}(s, \bar{m}(s), \bar{u}(s)) ds + K(T, \bar{m}(T))$, that is, the cost obtained by applying the nominal control input \bar{u} to the system starting from $\bar{m}(t)$ at time t . To explain the second-order term in the above expansion, by the definition, the second derivative $D^2 V(t, \bar{m})$ of the real-valued function $V(t, \cdot) : \mathcal{M} \rightarrow \mathbb{R}$ evaluated at \bar{m} is a bounded linear map from \mathcal{M} to \mathcal{M}^* (Lang, 1999). In our notation, $D^2 V(t, \bar{m}) \cdot \delta m$ denotes the evaluation of $D^2 V(t, \bar{m})$ at δm , giving an element in \mathcal{M}^* that can be paired with $\delta m \in \mathcal{M}$. Conceptually, $D^2 V(t, \bar{m})$ is nothing but the infinite-dimensional Hessian matrix with the (i, j) -entry given by $\frac{\partial^2 V}{\partial z_i \partial z_j} |_{(t, \bar{m})}$. If δm is small enough to ensure a sufficiently accurate approximation of the value function $V(t, \bar{m})$ up to the second-order terms, then we integrate (29) with the $o(\delta m^2)$ term neglected into the Hamilton-Jacobi-Bellman equation, yielding

$$\begin{aligned} \frac{\partial}{\partial t} \bar{V}(t, \bar{m}) + \frac{\partial}{\partial t} \delta V(t, \bar{m}) + \left\langle \frac{\partial}{\partial t} DV(t, \bar{m}), \delta m \right\rangle + \frac{1}{2} \left\langle \frac{\partial}{\partial t} D^2 V(t, \bar{m}) \cdot \delta m, \delta m \right\rangle \\ + \min_{\delta u} \left\{ H(t, \bar{m} + \delta m, \bar{u} + \delta u, DV(t, \bar{m})) + \left\langle D^2 V(t, \bar{m}) \cdot \delta m, \bar{F}(t, \bar{m} + \delta m, \bar{u} + \delta u) \right\rangle \right\} = 0 \end{aligned} \quad (30)$$

with the system Hamiltonian $H(t, m, u, DV) = \bar{r}(t, m, u) + \langle DV(t, m), \bar{F}(t, m, u) \rangle$, which is the key equation to the development of the algorithm for successively improving the nominal control policy \bar{u} .

Second-order policy search. To initialize the algorithm, we start with a nominal control policy u_0 , applying which to the moment system generates a nominal trajectory $m^{(0)}$. Because the pair $(m^{(0)}, u^{(0)})$ satisfies the system, the variation of the trajectory is $\delta m = 0$. In this case, the second-order expanded Hamilton-Jacobi-Bellman equation in (30) takes the form

$$\frac{\partial}{\partial t} \bar{V}(t, m^{(0)}) + \frac{\partial}{\partial t} \delta V(t, m^{(0)}) + \min_{\delta u} H(t, m^{(0)}, u^{(0)} + \delta u, DV(t, m^{(0)})) = 0, \quad (31)$$

and in many cases, the minimization of the Hamiltonian can be solved analytically, giving a new control policy $u^{(0)*}$, represented in terms of $m^{(0)}$ and $DV(t, m^{(0)})$ as a feedback control. However, steered by this policy, the system trajectory may not be the nominal one anymore, and hence a variation δm is introduced to the nominal trajectory as $m^{(0)} + \delta m$. Correspondingly, the second-order expanded Hamilton-Jacobi-Bellman equation becomes the one in (30) with \bar{u} replaced by $u^{(0)*}$, in which the minimization is taken over δu for the function

$$H(t, m^{(0)} + \delta m, u^{(0)*} + \delta u, DV(t, m^{(0)})) + \left\langle D^2 V(t, m^{(0)}) \cdot \delta m, \bar{F}(t, m^{(0)} + \delta m, u^{(0)*} + \delta u) \right\rangle.$$

We further expand this function around $(m^{(0)}, u^{(0)*})$ up to the second-order terms, yielding

$$\begin{aligned} H + \left\langle \frac{\partial H}{\partial u}, \delta u \right\rangle + \left\langle DH + D^2 V \cdot \bar{F}, \delta m \right\rangle + \left\langle \left(\frac{\partial DH}{\partial u} + \frac{\partial \bar{F}'}{\partial u} \cdot D^2 V \right) \cdot \delta m, \delta u \right\rangle \\ + \frac{1}{2} \left\langle \frac{\partial^2 H}{\partial u^2} \cdot \delta u, \delta u \right\rangle + \frac{1}{2} \left\langle \left(D^2 H + D \bar{F}' \cdot D^2 V + D^2 V' \cdot DF \right) \cdot \delta m, \delta m \right\rangle, \end{aligned} \quad (32)$$

where the terms involving H , \bar{F} , and V are evaluated at $(t, m^{(0)}, u^{(0)*}, DV(t, m^{(0)}))$, $(t, m^{(0)}, u^{(0)*})$, and $(t, m^{(0)})$, respectively, and ‘ \prime ’ denotes the dual of a linear operator. For example, because $\bar{F}(t, m^{(0)}, u^{(0)*}) \in \mathcal{M}$ by identifying the tangent space of \mathcal{M} at m^0 with \mathcal{M} itself, $D\bar{F}(t, m^{(0)}, u^{(0)*}) : \mathcal{M} \rightarrow \mathcal{M}$ is a linear map, then its dual operator is defined as a linear map $D\bar{F}'(t, m^{(0)}, u^{(0)*}) : \mathcal{M}^* \rightarrow \mathcal{M}^*$ satisfying $\langle L, D\bar{F}(t, m^{(0)}, u^{(0)*}) \cdot z \rangle = \langle D\bar{F}'(t, m^{(0)}, u^{(0)*}) \cdot L, z \rangle$ for any $z \in \mathcal{M}$ and $L \in \mathcal{M}^*$ (Yosida, 1980). Conceptually, dual operators are nothing but transpose matrices. Next, to minimize the function in (32), the necessary condition is the vanishing of its derivative with respect to δu , giving

$$\frac{\partial H}{\partial u} + \frac{\partial^2 H}{\partial u^2} \cdot \delta u + \left(\frac{\partial DH}{\partial u} + \frac{\partial \bar{F}'}{\partial u} \cdot D^2 V \right) \cdot \delta m = 0, \quad (33)$$

in which it is necessary that $\frac{\partial H}{\partial u}|_{(t, m^{(0)}, u^{(0)*}, DV(t, m^{(0)}))} = 0$ since u minimizes the Hamiltonian $H(t, m^{(0)}, u^{(0)*}, DV(t, m^{(0)}))$ by our choice.

Recall that our intention is to approximate the Hamilton-Jacobi-Bellman equation up to the second-order term in δm . Therefore, it is required that δm and δu are in the same order, meaning, they satisfy a linear relationship; otherwise, say δu is quadratic in δm , then the terms $\left\langle \left(\frac{\partial DH}{\partial u} + \frac{\partial \bar{F}'}{\partial u} \cdot D^2 V \right) \cdot \delta m, \delta u \right\rangle$ and $\frac{1}{2} \left\langle \frac{\partial^2 H}{\partial u^2} \cdot \delta u, \delta u \right\rangle$ in (32) are of orders higher than δm^2 . Formally, there is a linear map $A : \mathcal{M} \rightarrow \mathbb{R}^m$ such that $\delta u = A \cdot \delta m$. To find A , we replace δu by $A \cdot \delta m$ in the necessary optimality condition in (33), leading to

$$A = - \left(\frac{\partial^2 H}{\partial u^2} \right)^{-1} \cdot \left(\frac{\partial DH}{\partial u} + \frac{\partial \bar{F}'}{\partial u} \cdot D^2 V \right).$$

With this choice of A , the function in (32) becomes

$$H + \left\langle DH + D^2V \cdot \bar{F}, \delta m \right\rangle + \frac{1}{2} \left\langle \left(D^2H + D\bar{F}' \cdot D^2V + D^2V' \cdot DF - A' \cdot D^2H \cdot A \right) \cdot \delta m, \delta m \right\rangle$$

so that the second-order expansion of the Hamilton-Jacobi-Bellman equation in (30) takes the form

$$\begin{aligned} \frac{\partial}{\partial t} \bar{V} + \frac{\partial}{\partial t} \delta V + \left\langle \frac{\partial}{\partial t} DV, \delta m \right\rangle + \frac{1}{2} \left\langle \frac{\partial}{\partial t} D^2V(t, \bar{m}) \cdot \delta m, \delta m \right\rangle + H + \left\langle DH + D^2V \cdot \bar{F}, \delta m \right\rangle \\ + \frac{1}{2} \left\langle \left(D^2H + D\bar{F}' \cdot D^2V + D^2V' \cdot DF - A' \cdot \frac{\partial^2 H}{\partial u^2} \cdot A \right) \cdot \delta m, \delta m \right\rangle = 0, \end{aligned} \quad (34)$$

Because δm is arbitrary, the coefficient of each order of δm must be 0, which transforms (34) into a system of three partial differential equations

$$\begin{aligned} \frac{\partial}{\partial t} \bar{V} + \frac{\partial}{\partial t} \delta V - H &= 0 \\ \frac{\partial}{\partial t} DV - DH - D^2V \cdot \bar{F} &= 0 \\ \frac{\partial}{\partial t} D^2V - D^2H - D\bar{F}' \cdot D^2V - D^2V' \cdot DF + A' \cdot \frac{\partial^2 H}{\partial u^2} H \cdot A &= 0 \end{aligned}$$

with terms involving H , \bar{F} , and V are evaluated at $(t, m^{(0)}, u^{(0)*}, DV(t, m^{(0)}))$, $(t, m^{(0)}, u^{(0)*})$, and $(t, m^{(0)})$, respectively, as before. Integrating with the chain rule as

$$\begin{aligned} \frac{d}{dt} (\bar{V} + \delta V) &= \frac{\partial}{\partial t} \bar{V} + \frac{\partial}{\partial t} \delta V + \langle DV, \bar{F}(t, m^{(0)}, u^{(0)}) \rangle \\ \frac{d}{dt} DV &= \frac{\partial}{\partial t} DV + D^2V \cdot \bar{F}(t, m^{(0)}, u^{(0)}) \\ \frac{d}{dt} D^2V &= \frac{\partial}{\partial t} D^2V + \frac{1}{2} D^3V \cdot \bar{F}(t, m^{(0)}, u^{(0)}) + \frac{1}{2} \bar{F}(t, m^{(0)}, u^{(0)})' \cdot D^3V' \end{aligned}$$

then gives three ordinary differential equations

$$\frac{d}{dt} \delta V(t, m^{(0)}(t)) = H(t, m^{(0)}, u^{(0)}, DV(t, m^{(0)})) - H(t, m^{(0)}, u^{(0)*}, DV(t, m^{(0)})), \quad (35)$$

$$\frac{d}{dt} DV(t, m^{(0)}(t)) = -DH - D^2V \cdot (\bar{F}(t, m^{(0)}, u^{(0)*}) - \bar{F}(t, m^{(0)}, u^{(0)})), \quad (36)$$

$$\begin{aligned} \frac{d}{dt} D^2V(t, m^{(0)}(t)) &= -D^2H - D\bar{F}' \cdot D^2V - D^2V' \cdot D\bar{F} \\ &\quad + \left[\frac{\partial DH}{\partial u} + \frac{\partial \bar{F}'}{\partial u} \cdot D^2V \right]' \cdot \left(\frac{\partial^2 H}{\partial u^2} \right)^{-1} \cdot \left[\frac{\partial DH}{\partial u} + \frac{\partial \bar{F}'}{\partial u} \cdot D^2V \right], \end{aligned} \quad (37)$$

where we use $\frac{d}{dt} \bar{V}(t, m^{(0)}(t)) = \frac{d}{dt} \left\{ \int_t^T \bar{r}(s, m^{(0)}(s), u^{(0)}(s)) ds + \bar{K}(T, m^{(0)}(T)) \right\} = -\bar{r}(t, m^{(0)}(t), u^{(0)}(t))$ and $A = - \left(\frac{\partial^2 H}{\partial u^2} \right)^{-1} \cdot \left(\frac{\partial DH}{\partial u} + \frac{\partial \bar{F}'}{\partial u} \cdot D^2V \right)$, and omit the third-order

terms involving D^3V . Moreover, because the value function satisfies $V(T, m^{(0)}(T)) = \bar{K}(T, m^{(0)}(T))$, we have the terminal conditions for the ordinary differential equations in (35) to (36) as $\delta V(T, m^{(0)}(T)) = 0$, $DV(T, m^{(0)}(T)) = D\bar{K}(T, m^{(0)}(T))$, and $D^2V(T, m^{(0)}(T)) = D^2\bar{K}(T, m^{(0)}(T))$. In particular, the data $DV(t, m^{(0)}(t))$ obtained from solving the above systems of differential equations is then used to compute the control policy $u^{(0)*}$, which has been represented in terms of $\bar{m}(t)$ and $DV(t, m^{(0)}(t))$ when minimizing the Hamiltonian in (31) and hence gives rise to an improvement of the nominal control policy $u^{(0)}$. This in turn completes one iteration of the proposed policy search algorithm.

Appendix D. Derivation of Moment Systems

D.1 Infinite-dimensional LQR

In the following, we consider the finite-time horizon LQR problem

$$\begin{cases} \frac{d}{dt}x(t, \beta) = \beta x(t, \beta) + u(t), \\ V(t, x_t) = \int_{\Omega} \left[\int_0^T (x^2(t, \beta) + u^2(t)) dt d\beta + x^2(T, \beta) \right] d\beta, \end{cases} \quad (38)$$

where $\Omega = [-1, 1]$ and $x_t \in L^2(\Omega)$, the space of real-valued square-integrable functions defined on Ω .

Moment kernelization. We pick $\{\Phi_k\}_{k \in \mathbb{N}}$ to be the set of Chebyshev polynomials, and by using the recursive relation of Chebyshev polynomials $2\Phi_k = \Phi_{k-1} + \Phi_{k+1}$, we have

$$\begin{aligned} \frac{d}{dt}m_k(t) &= \frac{d}{dt} \int_{-1}^1 \Phi_k(\beta) x(t, \beta) d\beta = \int_{-1}^1 \Phi_k(\beta) \frac{d}{dt}x(t, \beta) d\beta = \int_{-1}^1 \Phi_k(\beta) [\beta x(t, \beta) + u(t)] d\beta \\ &= \frac{1}{2} \int_{-1}^1 [\Phi_{k-1}(\beta) + \Phi_{k+1}(\beta)] x(t, \beta) d\beta + \int_{-1}^1 \Phi_k(\beta) d\beta \cdot u(t) \\ &= \frac{1}{2} [m_{k-1}(t) + m_{k+1}(t)] + b_k u(t), \end{aligned}$$

where the change of the integral and time derivative follows from the dominated convergence theorem (Folland, 2013), Φ_{-1} , and hence m_{-1} , are defined to be identically 0, and b_k is given by $\frac{(-1)^k + 1}{1 - k^2}$ for $k \neq 1$ and 0 otherwise. We further let $L : \ell^2 \rightarrow \ell^2$ and $R : \ell^2 \rightarrow \ell^2$ denote the left and right shift operators, given by, $(m_0(t), m_1(t), \dots) \mapsto (m_1(t), m_2(t), \dots)$ and $(m_0(t), m_1(t), \dots) \mapsto (0, m_0(t), \dots)$, respectively, then the moment system associated with the linear ensemble system in (38) is a linear system evolving on ℓ^2 in the form $\frac{d}{dt}m(t) = Am(t) + Bu(t)$ with $A = \frac{1}{2}(L + R)$ and $B \in \ell^2$ whose k^{th} component is b_k . On the other hand, to parameterize the cost functional, we note that the moment transformation, that is, the Fourier transform, is a unitary operator from $L^2(\Omega)$ to ℓ^2 , as a result of which $J(u) = \int_0^1 [\|m(t)\|^2 + 2u^2(t)] dt + \|m(T)\|^2$ in the moment parameterization, where $\|\cdot\|$ denotes the ℓ^2 norm. In summary, the LQR problem in (38) in the moment kernel parameterization has the form

$$\begin{cases} \frac{d}{dt}m(t) = Am(t) + Bu(t), \\ V(t, m(T)) = \int_0^1 [\|m(t)\|^2 + 2u^2(t)] dt + \|m(T)\|^2. \end{cases} \quad (39)$$

Second-order policy search. The Hamiltonian $H : \ell^2 \times \mathbb{R} \times \ell^2 \rightarrow \mathbb{R}$ of the moment system in (39) is given by $(z, a, p) \mapsto \|z\|^2 + 2a^2 + \langle p, Az + Ba \rangle$, in which $\langle \cdot, \cdot \rangle$ is the ℓ^2 inner product. Let $V : [0, T] \times \ell^2 \rightarrow \mathbb{R}$ be the value function, then along a trajectory $m(t)$ of the system, by setting $\frac{\partial H}{\partial a}|_{(m(t), a, DV(t, m(t)))} = 4a + \langle DV(m(t)), B \rangle = 0$, we obtain $u^*(t) = -\frac{1}{4}\langle DV(t, m(t)), B \rangle$. Consequently, the differential equations in (35) to (37) for the policy improvement algorithm read

$$\begin{aligned} \frac{d}{dt}\delta V(t, m(t)) &= \left[u + \frac{1}{4}\langle DV, B \rangle \right]^2, \\ \frac{d}{dt}DV(t, m(t)) &= -2m - A' \cdot DV - D^2V \cdot \left(u - \frac{1}{4}\langle DV, B \rangle \right), \\ \frac{d}{dt}D^2V(t, m(t)) &= -2I - A' \cdot D^2V - D^2V \cdot A + \frac{1}{4}D^2V \cdot B \cdot B' \cdot D^2V, \end{aligned}$$

where I denotes the identity operator on ℓ^2 and we use the fact that D^2V is a self-adjoint operator on ℓ^2 . More concretely, when applying Algorithm 2 to a truncated moment system, say of truncation order N , then A and B in the above system of differential equations are replaced by

$$\hat{A}_N = \frac{1}{2} \begin{bmatrix} 0 & 1 & & & \\ 1 & 0 & 1 & & \\ & 1 & 0 & \ddots & \\ & & \ddots & \ddots & 1 \\ & & & 1 & 0 \end{bmatrix} \in \mathbb{R}^{(N+1) \times (N+1)} \quad \text{and} \quad \hat{B}_N = \begin{bmatrix} b_0 \\ b_1 \\ b_2 \\ \vdots \\ b_N \end{bmatrix} \in \mathbb{R}^{N+1},$$

respectively, and the operator dual is essentially the matrix transpose.

D.2 Moment kernelization of nuclear spin systems

The policy learning problem for the nuclear spin systems in (26) is given by

$$\begin{cases} \frac{d}{dt}x(t, \beta) = \beta[u(t)\Omega_y + v(t)\Omega_x]x(t, \beta), \\ V(t, x_t) = \int_t^T (u^2(t) + v^2(t))dt + \int_{1-\delta}^{1+\delta} |x(T, \beta) - x_F(\beta)|^2 d\beta, \end{cases}$$

where the Bloch system is defined on $L^2(\Omega, \mathbb{R}^3)$ with $\Omega = [1 - \delta, 1 + \delta]$ for some $0 < \delta < 1$.

Moment kernelization. Similar to the LQR case presented above, we still define the moments by using the set of Chebyshev polynomials $\{\Phi_k\}_{k \in \mathbb{N}}$. However, in order to fully utilize the orthonormal property of Chebyshev polynomials, which only holds on $\Omega' = [-1, 1]$, we consider the transformation $\psi : \Omega \rightarrow \Omega'$, given by $\beta \mapsto (\beta - 1)/\delta$, and defined the moments by

$$m_k(t) = \int_{\Omega} \Phi_k \circ \psi \cdot x_t d\lambda = \int_{\Omega'} \Phi \cdot x_t \circ \psi^{-1} d\psi_{\#}\lambda,$$

where λ denotes the Lebesgue measure on \mathbb{R} , and $\psi_{\#}\lambda$ is the pushforward measure of λ by ψ , that is, $\psi_{\#}\lambda(I) = \lambda(\psi^{-1}(I)) = \delta\lambda(I)$ is satisfied for $I \subseteq \mathbb{R}$, equivalently $d\psi_{\#}\lambda = \delta d\lambda$. This directly implies that the cost functional in the moment parameterization takes the

form $J(u, v) = \int_0^T (u^2(t) + v^2(t))dt + \|m(T)\|^2$, where $\|\cdot\|$ denotes the ℓ^2 -norm on the \mathbb{R}^3 -valued sequences, given by, $\|m(T)\|^2 = \sum_{k=0}^\infty |m_k(T)|^2$. Next, we compute the moment parameterization of the Bloch system in (26) as follows

$$\begin{aligned} \frac{d}{dt}m_k(t) &= \frac{d}{dt} \int_{-1}^1 \Phi_k(\eta) \cdot x(t, \psi^{-1}(\eta)) d\psi_{\#}\lambda(\eta) = \int_{-1}^1 \Phi_k(\eta) \frac{d}{dt}x(t, \psi^{-1}(\eta)) d\psi_{\#}\lambda(\eta) \\ &= \int_{-1}^1 \Phi_k(\eta) (\delta\eta + 1) [u(t)\Omega_y + v(t)\Omega_x] x(t, \psi^{-1}(\eta)) d\psi_{\#}\lambda(\eta) \\ &= [u(t)\Omega_y + v(t)\Omega_x] \cdot \left\{ \frac{\delta}{2} [m_{k-1}(t) + m_{k+1}(t)] + m_k(t) \right\}, \end{aligned}$$

where we use the recursive relation satisfied by Chebyshev polynomials, i.e., $2\Phi_k = \Phi_{k-1} + \Phi_{k+1}$, together with $\Phi_{-1} = 0$. Let L and R be the left and right shift operators for real-valued sequences as introduced in Section 4.2.1, the moment parameterization of the Bloch ensemble system is given by $\frac{d}{dt}m(t) = \left[\frac{\delta}{2}(R + L) + I \right] \otimes [u(t)\Omega_y + v(t)\Omega_x] m(t)$, where I denotes the identity operator for real-valued sequences and \otimes be the tensor product of linear operators. As a result, we have obtained the moment kernel parameterization of the RL problem for the Bloch system as

$$\begin{cases} \frac{d}{dt}m(t) = [u(t)B_y + v(t)B_x]m(t), \\ V(t, m(t)) = \int_0^T (u^2(t) + v^2(t))dt + \|m(T) - m_F\|^2 \end{cases} \quad (40)$$

where we define $B_x = [\delta(R + L)/2 + I] \otimes \Omega_x$ and $B_y = [\delta(R + L)/2 + I] \otimes \Omega_y$ to simplify the notations, and m_0 and m_F are the moment sequences of the constant functions $x_0(\beta) = (0, 0, 1)'$ and $x_F(\beta) = (1, 0, 0)'$, respectively.

Second-order policy search. The Hamiltonian of the moment system in (40) is given by $H : (\ell^2)^3 \times \mathbb{R} \times \mathbb{R} \times (\ell^2)^3 \rightarrow \mathbb{R}$, $(z, a, b, p) \mapsto a^2 + b^2 + \langle p, [aB_y + bB_x]z \rangle$. Let $V : [0, T] \times \ell^2 \rightarrow \mathbb{R}$ be the value function, then along a moment trajectory, by setting $\frac{\partial H}{\partial a}|_{(m(t), a, b, DV(t, m(t)))} = 2a + \langle DV, B_y m(t) \rangle = 0$ and $\frac{\partial H}{\partial b}|_{(m(t), a, b, DV(t, m(t)))} = 2b + \langle DV, B_x m(t) \rangle = 0$, we obtain the optimal policies $u^*(t) = -\langle DV, B_y m(t) \rangle / 2$ and $v^*(t) = -\langle DV, B_x m(t) \rangle / 2$. Integrating them into the system of differential equations in (35) to (37) for the policy improvement algorithm yields

$$\begin{aligned} \frac{d}{dt}\delta V(t, m(t)) &= (u - u^*)^2 + (v - v^*)^2, \\ \frac{d}{dt}DV(t, m(t)) &= -DV \cdot (u^*B_y + v^*B_x) - D^2V \cdot [(u^* - u)B_y + (v^* - v)B_x]m, \\ \frac{d}{dt}D^2V(t, m(t)) &= -(u^*B_y + v^*B_x)' \cdot D^2V - D^2V \cdot (u^*B_y + v^*B_x), \\ &\quad + \frac{1}{2}(DV \cdot B_y + m' \cdot B_y' \cdot D^2V)' \cdot (DV \cdot B_y + m' \cdot B_y' \cdot D^2V) \\ &\quad + \frac{1}{2}(DV \cdot B_x + m' \cdot B_x' \cdot D^2V)' \cdot (DV \cdot B_x + m' \cdot B_x' \cdot D^2V). \end{aligned}$$

Specifically, when applying Algorithm 2 to the order N truncated problem, B_x and B_y are replaced by the $\mathbb{R}^{3(N+1) \times 3(N+1)}$ block matrices

$$(\widehat{B}_x)_N = \frac{1}{2} \begin{bmatrix} 0 & \Omega_x & & & \\ \Omega_x & 0 & \Omega_x & & \\ & \Omega_x & 0 & \ddots & \\ & & \ddots & \ddots & \Omega_x \\ & & & \Omega_x & 0 \end{bmatrix} \quad \text{and} \quad (\widehat{B}_y)_N = \frac{1}{2} \begin{bmatrix} 0 & \Omega_y & & & \\ \Omega_y & 0 & \Omega_y & & \\ & \Omega_y & 0 & \ddots & \\ & & \ddots & \ddots & \Omega_y \\ & & & \Omega_y & 0 \end{bmatrix},$$

respectively.

References

- A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science*, 362(6419):1140–1144, 2018.
- S.V. Albrecht, F. Christianos, and L. Schäfer. *Multi-Agent Reinforcement Learning: Foundations and Modern Approaches*. MIT Press, 2024.
- Vladimir I. Arnold. *Ordinary Differential Equations*. MIT Press, 1978.
- Bowen Baker, Ingmar Kanitscheider, Todor Markov, Yi Wu, Glenn Powell, Bob McGrew, and Igor Mordatch. Emergent tool use from multi-agent autotutorials. In *International Conference on Learning Representations*, 2020.
- Aaron Becker and Timothy Bretl. Approximate steering of a unicycle under bounded model perturbation using ensemble control. *IEEE Transactions on Robotics*, 28(3):580–591, 2012.
- Marc G. Bellemare, Salvatore Candido, Pablo Samuel Castro, Jun Gong, Marlos C. Machado, Subhdeep Moitra, Sameera S. Ponda, and Ziyu Wang. Autonomous navigation of stratospheric balloons using reinforcement learning. *Nature*, 588(7836):77–82, 2020.
- R. Bellman. *Adaptive Control Processes: A Guided Tour*. Princeton Legacy Library. Princeton University Press, 1961.
- R. Bellman, Rand Corporation, and Karreman Mathematics Research Collection. *Dynamic Programming*. Rand Corporation research study. Princeton University Press, 1957.
- A. Bensoussan, J. Frehse, and P. Yam. *Mean Field Games and Mean Field Type Control Theory*. SpringerBriefs in Mathematics. Springer New York, 2013. ISBN 9781461485087.
- D. Bertsekas and J.N. Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, 1996.
- Dimitri Bertsekas. *Reinforcement Learning and Optimal Control*. Athena Scientific, 2019.
- P. Billingsley. *Probability and Measure*, volume 245 of *Wiley Series in Probability and Statistics*. Wiley, 3 edition, 1995.
- Roger W. Brockett. *Finite Dimensional Linear Systems*, volume 74 of *Classics in Applied Mathematics*. Society for Industrial and Applied Mathematics, 2015.

- Marin Bukov, Alexandre G. R. Day, Dries Sels, Phillip Weinberg, Anatoli Polkovnikov, and Pankaj Mehta. Reinforcement learning in different phases of quantum control. *Phys. Rev. X*, 8:031086, Sep 2018.
- Lucian Buşoniu, Robert Babuška, and Bart De Schutter. *Multi-agent Reinforcement Learning: An Overview*, pages 183–221. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010.
- René Carmona, Mathieu Laurière, and Zongjun Tan. Linear-quadratic mean-field reinforcement learning: convergence of policy gradient methods. *arXiv preprint arXiv:1910.04295*, 2019.
- René Carmona, Kenza Hamidouche, Mathieu Laurière, and Zongjun Tan. Policy optimization for linear-quadratic zero-sum mean-field type games. In *2020 59th IEEE Conference on Decision and Control (CDC)*, pages 1038–1043, 2020.
- John Cavanagh, Nicholas J. Skelton, Wayne J. Fairbrother, Mark Rance, and III Arthur G. Palmer. *Protein NMR Spectroscopy: Principles and Practice*. Elsevier, 2 edition, 2010.
- Chunlin Chen, Daoyi Dong, Ruixing Long, Ian R. Petersen, and Herschel A. Rabitz. Sampling-based learning control of inhomogeneous quantum ensembles. *Phys. Rev. A*, 89:023402, Feb 2014. doi: 10.1103/PhysRevA.89.023402.
- X. Chen, E. Torrontegui, D. Stefanatos, J.-S. Li, and J. G. Muga. Optimal trajectories for efficient atomic transport without final excitation. *Physical Review A*, 84:043415, 2011.
- ShiNung Ching and Jason T. Ritt. Control strategies for underactuated neural ensembles driven by optogenetic stimulation. *Front Neural Circuits*, 7:54, 2013.
- Daoyi Dong, Chunlin Chen, Hanxiong Li, and Tzyh-Jong Tarn. Quantum reinforcement learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 38(5):1207–1220, 2008.
- Ethan N. Evans, Marcus A. Periera, George I. Boutselis, and Evangelos A. Theodorou. Variational optimization based reinforcement learning for infinite dimensional stochastic systems. In Leslie Pack Kaelbling, Danica Kragic, and Komei Sugiura, editors, *Proceedings of the Conference on Robot Learning*, volume 100 of *Proceedings of Machine Learning Research*, pages 1231–1246. PMLR, 30 Oct–01 Nov 2020.
- Lawrence C. Evans. *Partial Differential Equations*, volume 19 of *Graduate Studies in Mathematics*. American Mathematical Society, 2nd edition, 2010.
- Jakob Foerster, Ioannis Alexandros Assael, Nando de Freitas, and Shimon Whiteson. Learning to communicate with deep multi-agent reinforcement learning. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.
- Gerald B. Folland. *Real Analysis: Modern Techniques and Their Applications*, volume 40 of *Pure and Applied Mathematics: A Wiley Series of Texts, Monographs and Tracts*. John Wiley & Sons, 2 edition, 2013.

- Vincent François-Lavet, Peter Henderson, Riashat Islam, Marc G. Bellemare, and Joelle Pineau. An introduction to deep reinforcement learning. *Foundations and Trends® in Machine Learning*, 11(3-4):219–354, 2018.
- Zuyue Fu, Zhuoran Yang, Yongxin Chen, and Zhaoran Wang. Actor-critic provably finds nash equilibria of linear-quadratic mean-field games. In *International Conference on Learning Representations*, 2020.
- Carles Gelada, Saurabh Kumar, Jacob Buckman, Ofir Nachum, and Marc G. Bellemare. DeepMDP: Learning continuous latent space models for representation learning. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2170–2179. PMLR, 09–15 Jun 2019.
- S. J. Glaser, T. Schulte-Herbrüggen, M. Sieveking, N. C. Nielsen O. Schedletzy, O. W. Sørensen, and C. Griesinger. Unitary control in quantum ensembles, maximizing signal intensity in coherent spectroscopy. *Science*, 280:421–424, 1998.
- Jayesh K. Gupta, Maxim Egorov, and Mykel Kochenderfer. Cooperative multi-agent control using deep reinforcement learning. In Gita Sukthankar and Juan A. Rodriguez-Aguilar, editors, *Autonomous Agents and Multiagent Systems*, pages 66–83, Cham, 2017. Springer International Publishing. ISBN 978-3-319-71682-4.
- Hans L. Hamburger. Über eine erweiterung des stieltjesschen momentenproblems. *Mathematische Annalen*, 82:120–164, 1920.
- Hans L. Hamburger. Über eine erweiterung des stieltjesschen momentenproblems. *Mathematische Annalen*, 82:168–187, 1921a.
- Hans L. Hamburger. Über eine erweiterung des stieltjesschen momentenproblems. *Mathematische Annalen*, 81:235–319, 1921b.
- T. Hastie, R. Tibshirani, and J.H. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer series in statistics. Springer, 2009.
- Tobias Haug, Rainer Dumke, Leong-Chuan Kwek, Christian Miniatura, and Luigi Amico. Machine-learning engineering of quantum currents. *Phys. Rev. Res.*, 3:013034, Jan 2021.
- Felix Hausdorff. Momentprobleme für ein endliches intervall. *Mathematische Zeitschrift*, 16(1):220–248, 1923.
- Suzana Herculano-Houzel. The remarkable, yet not extraordinary, human brain as a scaled-up primate brain and its associated cost. *Proceedings of the National Academy of Sciences*, 109(supplement_1):10661–10668, 2012.
- Paulo Heredia, Hasan Ghadialy, and Shaoshuai Mou. Finite-sample analysis of distributed q-learning for multi-agent networks. In *2020 American Control Conference (ACC)*, pages 3511–3516, 2020.

- Paulo Heredia, Jemin George, and Shaoshuai Mou. Distributed offline reinforcement learning. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 4621–4626, 2022.
- Paulo C. Heredia and Shaoshuai Mou. Distributed multi-agent reinforcement learning by actor-critic method. *IFAC-PapersOnLine*, 52(20):363–368, 2019. ISSN 2405-8963. 8th IFAC Workshop on Distributed Estimation and Control in Networked Systems NECSYS 2019.
- David H. Jacobson and David Q. Mayne. *Differential Dynamic Programming*. Modern Analytic and Computational Methods in Science and Mathematics. American Elsevier Publishing Company, 1970.
- Wei-Cheng Jiang, Vignesh Narayanan, and Jr-Shin Li. Model learning and knowledge sharing for cooperative multiagent systems in stochastic environment. *IEEE Transactions on Cybernetics*, 51(12):5717–5727, 2021.
- Lukasz Kaiser, Mohammad Babaeizadeh, Piotr Milos, Blazej Osinski, Roy H Campbell, Konrad Czechowski, Dumitru Erhan, Chelsea Finn, Piotr Kozakowski, Sergey Levine, Afroz Mohiuddin, Ryan Sepassi, George Tucker, and Henryk Michalewski. Model based reinforcement learning for atari. In *International Conference on Learning Representations*, 2020.
- Alec Koppel, Garrett Warnell, Ethan Stump, Peter Stone, and Alejandro Ribeiro. Policy evaluation in continuous mdp with efficient kernelized gradient temporal difference. *IEEE Transactions on Automatic Control*, 66(4):1856–1863, 2021.
- Huibert Kwakernaak and Raphael Sivan. *Linear Optimal Control Systems*. Wiley-Interscience, 1972.
- Lucas Lamata. Basic protocols in quantum reinforcement learning with superconducting circuits. *Scientific Reports*, 7(1):1609, 2017. doi: 10.1038/s41598-017-01711-6.
- Serge Lang. *Fundamentals of Differential Geometry*, volume 191 of *Graduate Texts in Mathematics*. Springer New York, NY, 1999.
- Michael Laskin, Aravind Srinivas, and Pieter Abbeel. CURL: Contrastive unsupervised representations for reinforcement learning. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 5639–5650. PMLR, 13–18 Jul 2020.
- Mathieu Laurière, Sarah Perrin, Sertan Girgin, Paul Muller, Ayush Jain, Théophile Cabannes, Georgios Piliouras, Julien P’erolat, Romuald Élie, Olivier Pietquin, and Matthieu Geist. Scalable deep reinforcement learning algorithms for mean field games. In *International Conference on Machine Learning*, 2022.
- Ngan Le, Vidhiwar Singh Rathour, Kashu Yamazaki, Khoa Luu, and Marios Savvides. Deep reinforcement learning in computer vision: a comprehensive survey. *Artificial Intelligence Review*, 55(4):2733–2819, 2022.

- John M. Lee. *Introduction to Smooth Manifolds*, volume 218 of *Graduate Texts in Mathematics*. Springer New York, NY, 2nd edition, 2012.
- Guy Lever and Ronnie Stafford. Modelling policies in mdps in reproducing kernel hilbert space. In Guy Lebanon and S. V. N. Vishwanathan, editors, *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, volume 38 of *Proceedings of Machine Learning Research*, pages 590–598, San Diego, California, USA, 09–12 May 2015. PMLR.
- Jr-Shin Li. Control of inhomogeneous ensembles, May 2006.
- Jr-Shin Li. Ensemble control of finite-dimensional time-varying linear systems. *IEEE Transactions on Automatic Control*, 56(2):345–357, 2011.
- Jr-Shin Li and Navin Khaneja. Ensemble control of bloch equations. *IEEE Transactions on Automatic Control*, 54(3):528–536, 2009.
- Jr-Shin Li, Justin Ruths, Tsyr-Yan Yu, Haribabu Arthanari, and Gerhard Wagner. Optimal pulse design in quantum control: A unified computational method. *Proceedings of the National Academy of Sciences*, 108(5):1879–1884, 2011.
- Jr-Shin Li, Isuru Dasanayake, and Justin Ruths. Control and synchronization of neuron ensembles. *IEEE Transactions on Automatic Control*, 58(8):1919–1930, 2013. doi: 10.1109/TAC.2013.2250112.
- Jr-Shin Li, Wei Zhang, and Yuan-Hung Kuan. Moment quantization of inhomogeneous spin ensembles. *Annual Reviews in Control*, 54:305–313, 2022. ISSN 1367-5788.
- Jr-Shin Li, Yuan-Hung Kuan, and Wei Zhang. Optimal quantum control using ensemble quantization. *SIAM Journal on Control and Optimization*, 63(1):S107–S127, 2025.
- Michael L. Littman. Markov games as a framework for multi-agent reinforcement learning. In William W. Cohen and Haym Hirsh, editors, *Machine Learning Proceedings 1994*, pages 157–163. Morgan Kaufmann, San Francisco (CA), 1994.
- Siqi Liu, Guy Lever, Zhe Wang, Josh Merel, S. M. Ali Eslami, Daniel Hennes, Wojciech M. Czarnecki, Yuval Tassa, Shayegan Omidshafiei, Abbas Abdolmaleki, Noah Y. Siegel, Leonard Hasenclever, Luke Marris, Saran Tunyasuvunakool, H. Francis Song, Markus Wulfmeier, Paul Muller, Tuomas Haarnoja, Brendan Tracey, Karl Tuyls, Thore Graepel, and Nicolas Heess. From motor control to team play in simulated humanoid football. *Science Robotics*, 7(69):eabo0235, 2022.
- Pinxin Long, Tingxiang Fan, Xinyi Liao, Wenxi Liu, Hao Zhang, and Jia Pan. Towards optimally decentralized multi-robot collision avoidance via deep reinforcement learning. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6252–6259, 2018.
- Zehui Lu, Tianyu Zhou, and Shaoshuai Mou. Real-time multi-robot mission planning in cluttered environment. *Robotics*, 13(3), 2024.

- George W. Mackey. Harmonic analysis as the exploitation of symmetry - a historical survey. *Bulletin (New Series) of the American Mathematical Society.*, 3(1):543–698, 1980.
- W.J. Marks. William j. marks. *Current Treatment Options in Neurology*, 7:237–243, 2005.
- Maja J. Matarić. Reinforcement learning in the multi-robot domain. *Autonomous Robots*, 4(1):73–83, 1997. doi: 10.1023/A:1008819414322.
- David Q. Mayne. A second-order gradient method for determining optimal trajectories of non-linear discrete-time systems. *International Journal of Control*, 3(1):85–95, 1966.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharmashan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- I. Momennejad, E. M. Russek, J. H. Cheong, M. M. Botvinick, N. D. Daw, and S. J. Gershman. The successor representation in human reinforcement learning. *Nature Human Behaviour*, 1(9):680–692, 2017. doi: 10.1038/s41562-017-0180-8. URL <https://doi.org/10.1038/s41562-017-0180-8>.
- James R. Munkres. *Topology*. Prentice Hall, Incorporated, 2000.
- Tenavi Nakamura-Zimmerer, Qi Gong, and Wei Kang. Adaptive deep learning for high-dimensional hamilton–jacobi–bellman equations. *SIAM Journal on Scientific Computing*, 43(2):A1221–A1247, 2021.
- Vignesh Narayanan, Wei Zhang, and Jr-Shin Li. Duality of ensemble systems through moment representations. *IEEE Transactions on Automatic Control*, pages 1–8, 2024.
- Katsuyuki Nishimura, Riqiang Fu, and Timothy A. Cross. The effect of rf inhomogeneity on heteronuclear dipolar recoupling in solid state nmr: Practical performance of sfam and redor. *Journal of Magnetic Resonance*, 152(2):227–233, 2001. ISSN 1090-7807.
- B. Øksendal. *Stochastic Differential Equations: An Introduction with Applications*. Universitext (1979). Springer Berlin, Heidelberg, 6 edition, 2003.
- OpenAI, :, Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemyslaw Debiak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Chris Hesse, Rafal Józefowicz, Scott Gray, Catherine Olsson, Jakub Pachocki, Michael Petrov, Henrique P. d. O. Pinto, Jonathan Raiman, Tim Salimans, Jeremy Schlatter, Jonas Schneider, Szymon Sidor, Ilya Sutskever, Jie Tang, Filip Wolski, and Susan Zhang. Dota 2 with large scale deep reinforcement learning, 2019. URL <https://arxiv.org/abs/1912.06680>.
- Barna Pásztor, Andreas Krause, and Ilija Bogunovic. Efficient model-based multi-agent mean-field reinforcement learning. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL <https://openreview.net/forum?id=gvcDSDYUZX>.
- R.K. Pathria and P.D. Beale. *Statistical Mechanics*. Elsevier Science, 2021.

- V.I. Paulsen and M. Raghupathi. *An Introduction to the Theory of Reproducing Kernel Hilbert Spaces*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2016.
- I. Roos and K. Moelmer. Quantum computing with an inhomogeneously broadened ensemble of ions: Suppression of errors from detuning variations by specially adapted pulses and coherent population trapping. *Physical Review A*, 69:022321, 2004.
- Walter Rudin. *Principles of Mathematical Analysis*. McGraw-Hill, 3 edition, 1976.
- Nima Sarang and Charalambos Poullis. Tractable large-scale deep reinforcement learning. *Computer Vision and Image Understanding*, 232:103689, 2023. ISSN 1077-3142. doi: <https://doi.org/10.1016/j.cviu.2023.103689>. URL <https://www.sciencedirect.com/science/article/pii/S1077314223000693>.
- Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepel, Timothy Lillicrap, and David Silver. Mastering atari, go, chess and shogi by planning with a learned model. *Nature*, 588(7839):604–609, 2020.
- John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 1889–1897, Lille, France, 07–09 Jul 2015. PMLR.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms, 2017.
- Shiva Shahrokhi, Lillian Lin, Chris Ertel, Mable Wan, and Aaron T. Becker. Steering a swarm of particles using global inputs and swarm statistics. *IEEE Transactions on Robotics*, 34(1):207–219, 2018.
- Ali H. Shoeb. Application of machine learning to epileptic seizure onset detection and treatment, September 2009.
- David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel, and Demis Hassabis. Mastering the game of go without human knowledge. *Nature*, 550(7676):354–359, 2017. doi: 10.1038/nature24270.
- M. S. Silver, R. I. Joseph, and D. I. Hoult. Selective spin inversion in nuclear magnetic resonance and coherent optics through an exact solution of the bloch-riccati equation. *Physical Review A*, 31(4):2753–2755, 1985.
- Eduardo D. Sontag. *Mathematical Control Theory: Deterministic Finite Dimensional Systems*. Springer New York, NY, 2nd edition, 1998.
- D. Stefanatos and J.-S. Li. Minimum-time frictionless atom cooling in harmonic traps. *SIAM Journal on Control and Optimization*, 49:2440–2462, 2011.

- D. Stefanatos and J.-S. Li. Minimum-time quantum transport with bounded trap velocity. *IEEE Transactions on Automatic Control*, 59(3):733–738, 2014.
- Thomas J. Stieltjes. *Œuvres Complètes II - Collected Papers II*. Springer Collected Works in Mathematics. Springer Berlin, Heidelberg, 1993.
- Jayakumar Subramanian and Aditya Mahajan. Reinforcement learning in stationary mean-field games. Richland, SC, 2019. International Foundation for Autonomous Agents and Multiagent Systems.
- Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning An Introduction*. MIT Press, 2018.
- Oriol Vinyals, Igor Babuschkin, Wojciech M. Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H. Choi, Richard Powell, Timo Ewalds, Petko Georgiev, Junhyuk Oh, Dan Horgan, Manuel Kroiss, Ivo Danihelka, Aja Huang, Laurent Sifre, Trevor Cai, John P. Agapiou, Max Jaderberg, Alexander S. Vezhnevets, Rémi Leblond, Tobias Pohlen, Valentin Dalibard, David Budden, Yury Sulsky, James Molloy, Tom L. Paine, Caglar Gulcehre, Ziyu Wang, Tobias Pfaff, Yuhuai Wu, Roman Ring, Dani Yogatama, Dario Wünsch, Katrina McKinney, Oliver Smith, Tom Schaul, Timothy Lillicrap, Koray Kavukcuoglu, Demis Hassabis, Chris Apps, and David Silver. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019. doi: 10.1038/s41586-019-1724-z.
- Minh Vu, Bharat Singhal, Shen Zeng, and Jr-Shin Li. Data-driven control of oscillator networks with population-level measurement. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 34(3):033138, 03 2024.
- Scott B. Wilson. A neural network method for automatic and incremental learning applied to patient-dependent seizure detection. *Clinical Neurophysiology*, 116(8):1785–1795, August 2005.
- Yijing Xie, Shaoshuai Mou, and Shreyas Sundaram. Communication-efficient and resilient distributed q-learning. *IEEE Transactions on Neural Networks and Learning Systems*, 35(3):3351–3364, 2024.
- Erfu Yang and Dongbing Gu. Multiagent reinforcement learning for multi-robot systems: A survey. Technical report, tech. rep, 2004.
- Lin F. Yang and Mengdi Wang. Sample-optimal parametric q-learning using linearly additive features. In *International Conference on Machine Learning*, 2019.
- Yang Yang, Li Juntao, and Peng Lingling. Multi-robot path planning based on a deep reinforcement learning dqn algorithm. *CAAI Transactions on Intelligence Technology*, 5(3):177–183, 2020a.
- Yaodong Yang, Rui Luo, Minne Li, Ming Zhou, Weinan Zhang, and Jun Wang. Mean field multi-agent reinforcement learning. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of*

- Machine Learning Research*, pages 5567–5576, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018. PMLR.
- Zhuoran Yang, Chi Jin, Zhaoran Wang, Mengdi Wang, and Michael I. Jordan. On function approximation in reinforcement learning: optimism in the face of large state spaces. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS ’20, Red Hook, NY, USA, 2020b. Curran Associates Inc.
- Amir Yazdanbakhsh, Junchao Chen, and Yu Zheng. Menger: Massively large-scale distributed reinforcement learning. *NeurIPS, Beyond Backpropagation Workshop, 2020*, 2020. URL <https://beyondbackprop.github.io/>.
- Kōsaku Yosida. *Functional Analysis*, volume 123 of *Grundlehren der mathematischen Wissenschaften*. Springer Berlin, Heidelberg, 6 edition, 1980.
- Yao-Chi Yu, Wei Zhang, David O’Gara, Jr-Shin Li, and Su-Hsin Chang. A moment kernel machine for clinical data mining to inform medical decision making. *Scientific Reports*, 13(1):10459, 2023.
- Amy Zhang, Rowan Thomas McAllister, Roberto Calandra, Yarin Gal, and Sergey Levine. Learning invariant representations for reinforcement learning without reconstruction. In *International Conference on Learning Representations*, 2021a.
- Kaiqing Zhang, Zhuoran Yang, and Tamer Başar. *Multi-Agent Reinforcement Learning: A Selective Overview of Theories and Algorithms*. Springer International Publishing, Cham, 2021b.
- Wei Zhang and Jr-Shin Li. Uniform and selective excitations of spin ensembles with rf inhomogeneity. In *2015 54th IEEE Conference on Decision and Control (CDC)*, pages 5766–5771, 2015. doi: 10.1109/CDC.2015.7403125.
- A. Zlotnik and J.-S. Li. Optimal entrainment of neural oscillator ensembles. *Journal of Neural Engineering*, 9(4):046015, July 2012.