

A fast iterative PDE-based algorithm for feedback controls of nonsmooth mean-field control problems

Christoph Reisinger* Wolfgang Stockinger* Yufei Zhang*

Abstract. We propose a PDE-based accelerated gradient algorithm for optimal feedback controls of McKean–Vlasov dynamics that involve mean-field interactions both in the state and action. The method exploits a forward-backward splitting approach and iteratively refines the approximate controls based on the gradients of smooth costs, the proximal maps of nonsmooth costs, and dynamically updated momentum parameters. At each step, the state dynamics is approximated via a particle system, and the required gradient is evaluated through a coupled system of nonlocal linear PDEs. The latter is solved by finite difference approximation or neural network-based residual approximation, depending on the state dimension. We present exhaustive numerical experiments for low and high-dimensional mean-field control problems, including sparse stabilization of stochastic Cucker–Smale models, which reveal that our algorithm captures important structures of the optimal feedback control and achieves a robust performance with respect to parameter perturbation.

Key words. Controlled McKean–Vlasov diffusion, optimal gradient method, monotone scheme, neural network, sparse control, stochastic Cucker–Smale model.

AMS subject classifications. 49N80, 60H35, 35Q93, 93A16

1 Introduction

In this article, we propose a class of iterative methods for solving mean-field control (MFC) problems, where the state dynamics and cost functions depend upon the joint law of the state and the control processes. Let $T > 0$ be a given terminal time, $W = (W_t)_{t \in [0, T]}$ be an n -dimensional Brownian motion defined on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$, $\mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}$ the natural filtration of W augmented with an independent σ -algebra \mathcal{F}_0 , and $\mathcal{H}^2(\mathbb{R}^k)$ be the set of admissible controls containing all \mathbb{R}^k -valued square integrable \mathbb{F} -progressively measurable processes. For a given \mathcal{F}_0 -measurable initial state $\xi \in L^2(\Omega; \mathbb{R}^d)$ and a control $\alpha \in \mathcal{H}^2(\mathbb{R}^k)$, we consider the state process governed by the following controlled McKean–Vlasov diffusion:

$$dX_t = b(t, X_t, \alpha_t, \mathcal{L}_{(X_t, \alpha_t)}) dt + \sigma(t, X_t, \alpha_t, \mathcal{L}_{(X_t, \alpha_t)}) dW_t, \quad t \in [0, T]; \quad X_0 = \xi, \quad (1.1)$$

where $b : [0, T] \times \mathbb{R}^d \times \mathbb{R}^k \times \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^k) \rightarrow \mathbb{R}^d$ and $\sigma : [0, T] \times \mathbb{R}^d \times \mathbb{R}^k \times \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^k) \rightarrow \mathbb{R}^{d \times n}$ are sufficiently regular functions such that (1.1) admits a unique square integrable solution X^α . The value function of the optimal control problem is defined by

$$J^*(\xi) = \inf_{\alpha \in \mathcal{H}^2(\mathbb{R}^k)} J(\alpha; \xi), \quad \text{with } J(\alpha; \xi) = \mathbb{E} \left[\int_0^T \left(f(t, X_t^\alpha, \alpha_t, \mathcal{L}_{(X_t^\alpha, \alpha_t)}) + \ell(\alpha_t) \right) dt + g(X_T^\alpha, \mathcal{L}_{X_T^\alpha}) \right], \quad (1.2)$$

where $f : [0, T] \times \mathbb{R}^d \times \mathbb{R}^k \times \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^k) \rightarrow \mathbb{R}$, $g : \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$ are differentiable functions of at most quadratic growth, and $\ell : \mathbb{R}^k \rightarrow \mathbb{R} \cup \{\infty\}$ is a proper, lower semicontinuous and convex function. Above and hereafter, \mathcal{L}_U denotes the law of a random variable U , and $\mathcal{P}_2(E)$ denotes the Wasserstein space of probability measures on the Euclidean space E with finite second moment.

The above MFC problem extends classical stochastic control problems by allowing the mean-field interactions through the joint distribution of the state and control processes. It describes large population

*Mathematical Institute, University of Oxford, Oxford OX2 6GG, UK (christoph.reisinger@maths.ox.ac.uk, wolfgang.stockinger@maths.ox.ac.uk, yufei.zhang@maths.ox.ac.uk)

equilibria of interacting individuals controlled by a central planner, and plays an important role in economics [12, 1, 4], production management [31], biology [19, 9, 5] and social interactions [2, 3, 4]. Moreover, the extended real-valued function ℓ in (1.2) includes important examples such as characteristic functions of convex sets representing control constraints [1, 47], ℓ_1 -norm based regularizations used to induce sparsity or switching properties of the optimal control [19, 9], and entropy regularizations in machine learning [53, 33].

To solve (1.1)-(1.2), we aim to construct an optimal (decentralized) feedback control, i.e., a sufficiently regular function $\phi^* : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^k$ such that the corresponding controlled dynamics (1.1) (with $\alpha_t = \phi^*(t, X_t)$) admits a unique solution X^{ϕ^*} and $J^*(\xi) = J(\phi^*(\cdot, X^{\phi^*}), \xi)$.¹ The existence of such feedback controls has been shown in [47] (see also [20, Section 6.4] for special cases without mean-field interactions through control variables). The main advantage of a feedback strategy is that implementing the optimal control reduces to simple function evaluation at the current state of the system. Moreover, a feedback control allows us to interpret the mechanism of the optimal control. This is particularly important for economics and social science [41], where one would like to understand the cause of a decision, and explain its dependence on the state dynamics and the objective function. A feedback control also allows us to analyze the failure of certain control policies for fault diagnosis.

Existing numerical methods for MFC problems and their limitations. As analytical solutions to optimal feedback controls of (1.1)-(1.2) are rarely available, numerical schemes for solving such control problems become vital. Due to the nonlinear dependence on the marginal laws, it is difficult to follow the classical dynamic programming (DP) approach for constructing optimal feedback controls. The main challenge here is that to construct feedback controls, DP requires us to find derivatives of solutions to an infinite-dimensional Hamilton-Jacobi-Bellman (HJB) partial differential equation (PDE) defined on $[0, T] \times \mathcal{P}_2(\mathbb{R}^d)$ (see e.g., [46]), which is computationally intractable. Hence, most existing numerical methods are based on the optimization interpretation of (1.1)-(1.2).

The most straightforward approach for solving (1.1)-(1.2) is to restrict the optimization over feedback controls within a prescribed parametric family, i.e., the so-called policy gradient (PG) method (see e.g., [21]). It approximates the optimal feedback control in a parametric form depending on weights θ (for instance a deep neural network), and then seeks the optimal approximation by performing gradient descent of J with respect to the weights θ based on simulated trajectories of the state process. By exploiting an efficient neural network representation of the feedback control, the PG method can be adapted to solve MFC problems with high-dimensional state processes.

However, the PG method has two serious drawbacks, especially for solving control problems with nonlinear dynamics and nonsmooth costs in (1.1)-(1.2). Firstly, as the loss functional J is nonconvex and nonsmooth in the weights θ of the numerical feedback controls, there is no theoretical guarantee on the convergence of the PG method for solving nonsmooth MFC problems. In practice, even though the PG method may minimize the loss function reasonably well, the resulting approximate feedback control often fails to capture important structural features of the optimal feedback control, and consequently lacks the capability to provide sufficient insights of the optimal decision; see Figures 6, 7 and 9 in Section 5.2.2 where the PG method ignores the temporal and spatial nonlinearity and the sparsity of the optimal controls. Secondly, as the PG method computes approximate feedback controls based purely on sample trajectories of the state dynamics, it cannot recover optimal feedback controls outside the support of the optimal state process. Consequently, the performance of the approximate feedback control in general can be very sensitive to perturbations of the (random) initial state ξ of (1.1); see Section 5.1 for details. This is undesirable for practical applications of MFC problems, as the initial condition ξ describes the asymptotic regime of initial conditions of a large number of players, and often cannot be observed exactly.

Another approach is to solve the optimality systems arising from applying the Pontryagin Maximum Principle (PMP) to (1.1)-(1.2). Existing works consider special cases with neither nonsmooth costs nor mean-field interactions through control variables, and design numerical methods based on either a probabilistic or deterministic formulation (see e.g., [20, Section 6.2.4]). The probabilistic formulation represents an optimal control $\alpha^* \in \mathcal{H}^2(\mathbb{R}^k)$ as $\alpha_t^* = \hat{\alpha}(t, X_t^{\alpha^*}, \mathcal{L}_{X_t^{\alpha^*}}, Y_t^{\alpha^*}, Z^{\alpha^*})$ for $dt \otimes d\mathbb{P}$ -a.e., with $\hat{\alpha}$ being the pointwise minimizer of an associated Hamiltonian \mathcal{H} , and $(X^{\alpha^*}, Y^{\alpha^*}, Z^{\alpha^*})$ being the solution to a coupled forward-backward stochastic differential equation (FBSDE) depending on $\hat{\alpha}$. The coupled

¹ ϕ^* is called a decentralized control as it acts explicitly only on the time and state variables, while the dependence on the (deterministic) marginal laws of the optimal state and control processes is implicit through the time dependence.

FBSDE can be solved by first representing the solution in terms of grid functions [6, 48], binomial trees [6] or neural networks [26, 21, 28], and then employing regression methods to obtain the optimal approximation. Similarly, the deterministic formulation represents an optimal feedback control ϕ^* as $\phi^*(t, x) = \bar{\alpha}(t, x, \mu_t^*, (\nabla_x v)(t, x), (\text{Hess}_x v)(t, x))$ for all $(t, x) \in [0, T] \times \mathbb{R}^d$. Here $\bar{\alpha}$ is the pointwise minimizer of an associated Hamiltonian \mathcal{H} (possibly different from \mathcal{H}), and (μ^*, v) satisfy a coupled Fokker-Planck (FP)-HJB PDE system depending on $\bar{\alpha}$, which consists of a nonlinear FP equation for the marginal distribution μ^* of the optimal state process, and of a nonlinear HJB equation for the adjoint variable v . The FP-HJB system can then be solved by finite difference methods as in [2, 3, 4] or by neural network methods as in [22].

We observe, however, that the above PMP approach suffers from the following limitations. Firstly, the derivation of the optimality systems relies heavily on the analytic expression of pointwise minimizers of the Hamiltonians, which may not be available for general control problems. More crucially, as pointed out in [1], when there is a nonlinear dependence on the law of the control, the PMP in general cannot be expressed in terms of a pointwise minimization of Hamiltonian (see [47] for a detailed investigation of this issue). These factors prevent us from applying the PMP approach to solve (1.1)-(1.2) with general cost functions and control interactions. Secondly, similar to the PG method, solving the coupled FBSDE via regression focuses mainly along trajectories of the optimal state, and consequently would result in an approximate feedback control that is sensitive to perturbation of the initial state (see Section 5.1). Finally, solutions to the nonlinear FP equation in general only exist in the sense of distributions and often admit temporal and spatial singularity. This creates significant numerical challenges, especially when the diffusion coefficient σ of (1.1) is degenerate (as in most kinetic models) or the initial state ξ has a singular density; see Section 5 for concrete examples. In particular, in the high-dimensional setting, one may need to employ neural networks with complex structures to approximate such irregular solutions, which subsequently results in challenging optimization problems for training the networks.

Our contributions and related works. This paper proposes a class of iterative algorithms to construct optimal feedback controls for nonsmooth MFC problems (1.1)-(1.2).

- We construct a sequence of feedback controls $\phi^m : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^k$ whose realized control processes minimize the functional J . In the present setting, directly applying gradient descent to either the stochastic formulation of the control problem over open-loop controls, or to the deterministic reformulation over feedback controls has several critical deficiencies (see Sections 2.1 and 2.2 for details). To overcome these shortcomings, we propose a heuristic combination of the two formulations, and tailor Nesterov’s accelerated proximal gradient (NAG) method (also known as the Fast Iterative Shrinkage-Thresholding Algorithm [15]) to our problem. At each iteration with given feedback control, we evaluate the gradient of the smooth costs at the corresponding realized control process, and update the feedback control by incorporating the gradient information, the proximal map of the nonsmooth cost ℓ , and an explicit dynamically updated momentum parameter.

The proposed accelerated proximal gradient approach has the following advantages in solving (1.1)-(1.2): (i) unlike the aforementioned PMP approach, our algorithm requires neither (pointwise) analytical minimization of the Hamiltonian nor deriving the optimality systems, and hence can be applied to MFC problems with general mean-field interactions through the control variables; (ii) our algorithm shares the same computational complexity as the gradient-based algorithms in [45, 7, 53, 35], but enjoys an accelerated convergence rate and can handle general convex nonsmooth costs, including ℓ_1 -regularizers and control constraints. In fact, such an accelerated gradient iteration is known to be an optimal first order (gradient) method (in the sense of complexity analysis) for minimizing finite-dimensional nonsmooth functions [15]; (iii) our method represents the control iterates in a feedback form (cf., [7, 53, 35] which update controls as stochastic processes), and avoids the curse of dimensionality in the gradient evaluation as the number of iterations tends to infinity (see Section 2 for details); (iv) compared to directly applying NAG method to optimize (1.2) over feedback controls, our algorithm avoids the necessity of evaluating the proximal map of the nonsmooth cost with respect to feedback controls. This optimization enhances the efficiency of the algorithm; see Remark 2.2 for further details.

- We present a practical implementation of the above accelerated gradient algorithm by combining Monte Carlo and PDE approaches. At each iteration, the state dynamics with a given feedback

control is realized by using a particle approximation and Euler–Maruyama timestepping scheme, and the required gradient is computed by solving a coupled system of nonlocal linear PDEs, whose coefficients depend on the empirical measure of the particle system. The coupled PDE system is then solved with two different approaches depending on the state dimension d in (1.1), in order to balance the efficiency and computation complexity. In the low-dimensional setting (say $d \leq 2$), we discretize the coupled PDE system by a class of semi-implicit monotone finite difference approximations. To accommodate the curse of dimensionality with large state dimension, we also propose a residual approximation approach to solve the coupled PDE system, in which the numerical solution is decomposed into a pre-determined candidate solution and an unknown residual term. The computation of the residual term is addressed by a mesh-free method based on neural network approximation and stochastic optimization algorithms, such as the Stochastic Gradient Descent (SGD) algorithm or its variants.

The proposed algorithm combines the advantages of probabilistic and deterministic approaches. Firstly, the particle approximation allows for efficient computation of the marginal distribution of the state process and avoids the numerical challenge in solving a nonlinear FP equation (cf., [2, 3, 4, 22]). This is particularly relevant for high-dimensional MFC problems with degenerate diffusion coefficients or irregular initial distribution (see Section 5). Secondly, by exploiting the PDE formulation of the gradient evaluation, our algorithm recovers the optimal feedback control on the entire computational domain, rather than merely along the trajectories of the optimal state (cf., the probabilistic methods in [6, 21, 28, 48]). This allows us to capture important structures of the optimal control and achieve a robust performance with parameter uncertainty (see Section 5). As alluded to earlier, such an accurate approximation of the optimal feedback control is practically important for mathematical modelling and fault diagnosis in engineering. Finally, instead of directly applying SGD to high-dimensional PDE systems as in [52, 22], the proposed residual approximation approach leverages available efficient solvers to compute the dominant part of solutions (see e.g., the Riccati-based solvers in [31, 5]), and employs a small number of SGD iterations to fit the residual term. This significantly accelerates the convergence of the algorithm for solving high-dimensional MFC problems (see Figure 8).

- We demonstrate the effectiveness of the algorithm through extensive numerical experiments. This includes a two-dimensional nonsmooth MFC problem arising from portfolio liquidation with trade crowding, and a six-dimensional nonsmooth nonconvex MFC problem arising from sparse consensus control of stochastic Cucker–Smale models. Our experiments show that the resulting approximate feedback control correctly captures the temporal/spatial nonlinearity and the sparsity of the optimal control, and achieves a robust performance in the presence of initial state perturbation.

The rest of the paper is organized as follows. Section 2 describes our numerical methodology, including the accelerated proximal gradient iteration, the particle system for the state process, and the PDE system for the gradient evaluation. We then propose a class of finite difference approximations in Section 3 and neural network-based residual approximations in Section 4 to solve the PDE systems. In Section 5, we present exhaustive numerical experiments for multidimensional nonsmooth nonconvex MFC problems, which demonstrate that the proposed algorithm leads to more accurate and stable feedback controls than the aforementioned PG method and the PMP method.

2 Fast iterative PDE-based method for MFC problems

This section proposes a class of Markovian accelerated proximal gradient methods for MFC problems with nonsmooth running costs.

In the following subsections, we will incrementally introduce the main algorithm of the paper, in order to highlight the main ingredients and to stress which components (at least in special cases) are supported by theoretical results and which (at the present stage) are heuristic.² In Section 2.1, we will first formulate a gradient method with open-loop controls, without acceleration step and nonsmooth cost functional ℓ . Section 2.2 is dedicated to the formulation of a gradient method when the optimization is carried out over feedback functions.

²We thank an anonymous referee for their comments that helped bring clarity to the presentation.

These two different approaches serve as motivation for Section 2.3, which introduces a heuristic iterative PDE based algorithm to overcome certain deficiencies encountered in the first two formulations (both earlier sections assume $\ell \equiv 0$). Our main algorithm, which allows to incorporate nonsmooth cost and performs an additional moment step, is detailed in Section 2.4. Finally, Section 2.5 discusses existing convergence results for these various gradient algorithms.

2.1 First motivating algorithm: gradient descent in open-loop control

This section considers MFC problems with regular running costs ($\ell \equiv 0$ in (1.2)), and derives plain gradient-descent iterations over open-loop controls.

If $\ell \equiv 0$, (1.1)-(1.2) is a minimisation problem of $F : \mathcal{H}^2(\mathbb{R}^k) \rightarrow \mathbb{R}$ over the Hilbert space $\mathcal{H}^2(\mathbb{R}^k)$:

$$\inf_{\alpha \in \mathcal{H}^2(\mathbb{R}^k)} F(\alpha), \quad \text{with } F(\alpha) := \mathbb{E} \left[\int_0^T f(t, X_t^\alpha, \alpha_t, \mathcal{L}_{(X_t^\alpha, \alpha_t)}) dt + g(X_T^\alpha, \mathcal{L}_{X_T^\alpha}) \right], \quad (2.1)$$

where $X^\alpha \in \mathcal{S}^2(\mathbb{R}^d)$ is the state process controlled by α satisfying (1.1). By [1, Lemma 3.1] and by assuming differentiability of (b, σ, f) , F is Fréchet differentiable with derivative $\nabla F : \mathcal{H}^2(\mathbb{R}^k) \rightarrow \mathcal{H}^2(\mathbb{R}^k)$ satisfying for all $\alpha \in \mathcal{H}^2(\mathbb{R}^k)$,

$$(\nabla F)(\alpha)_t = (\partial_a H)(t, X_t^\alpha, \alpha_t, \mathcal{L}_{(X_t^\alpha, \alpha_t)}, Y_t^\alpha, Z_t^\alpha) + \tilde{\mathbb{E}}[(\partial_\nu H)(t, \tilde{X}_t^\alpha, \tilde{\alpha}_t, \mathcal{L}_{(X_t^\alpha, \alpha_t)}, \tilde{Y}_t^\alpha, \tilde{Z}_t^\alpha)(X_t^\alpha, \alpha_t)], \quad (2.2)$$

where $H : [0, T] \times \mathbb{R}^d \times \mathbb{R}^k \times \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^k) \times \mathbb{R}^d \times \mathbb{R}^{d \times n} \rightarrow \mathbb{R}$ is the Hamiltonian defined by:

$$H(t, x, a, \eta, y, z) := \langle b(t, x, a, \eta), y \rangle + \langle \sigma(t, x, a, \eta), z \rangle + f(t, x, a, \eta), \quad (2.3)$$

and (Y^α, Z^α) are square integrable adapted adjoint processes such that for all $t \in [0, T]$,

$$\begin{aligned} dY_t^\alpha &= - \left((\partial_x H)(t, X_t^\alpha, \alpha_t, \mathcal{L}_{(X_t^\alpha, \alpha_t)}, Y_t^\alpha, Z_t^\alpha) \right. \\ &\quad \left. + \tilde{\mathbb{E}}[(\partial_\mu H)(t, \tilde{X}_t^\alpha, \tilde{\alpha}_t, \mathcal{L}_{(X_t^\alpha, \alpha_t)}, \tilde{Y}_t^\alpha, \tilde{Z}_t^\alpha)(X_t^\alpha, \alpha_t)] \right) dt + Z_t^\alpha dW_t, \\ Y_T^\alpha &= (\partial_x g)(X_T^\alpha, \mathcal{L}_{X_T^\alpha}) + \tilde{\mathbb{E}}[(\partial_\mu g)(\tilde{X}_T^\alpha, \mathcal{L}_{X_T^\alpha})(X_T^\alpha)]. \end{aligned} \quad (2.4)$$

Above and hereafter, following [1], we use the tilde notation to denote an independent copy of a random variable. Moreover, for a given function $h : \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^k) \rightarrow \mathbb{R}$ and a measure $\eta \in \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^k)$ with marginals $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, $\nu \in \mathcal{P}_2(\mathbb{R}^k)$, we denote by $((\partial_\mu h)(\eta), (\partial_\nu h)(\eta))(\cdot) : \mathbb{R}^n \times \mathbb{R}^k \rightarrow \mathbb{R}^n \times \mathbb{R}^k$ the partial L-derivatives of h with respect to the marginals; see, e.g., [1, 20] for detailed definitions.

Based on the above interpretation, a gradient descent algorithm for optimal controls of (2.1) is given as follows: Let $\alpha^0 \in \mathcal{H}^2(\mathbb{R}^k)$ be the initial guess of the optimal control, and $\tau > 0$ be a chosen stepsize. Consider the sequence $(\alpha^m)_{m \in \mathbb{N}} \subset \mathcal{H}^2(\mathbb{R}^k)$ such that

$$\alpha_t^{m+1} = \alpha_t^m - \tau(\nabla F)(\alpha^m)_t, \quad \forall m \in \mathbb{N} \cup \{0\}, \quad (2.5)$$

dt \otimes dP-a.e., where $(\nabla F)(\alpha^m)$ is defined by (2.2).

Pros and cons of Iteration (2.5). The main advantage of (2.5) is that its convergence can be ensured given sufficient regularity of the functional $F : \mathcal{H}^2(\mathbb{R}^k) \rightarrow \mathbb{R}$ in (2.1). Indeed, as a special case of Theorem 3.1 of [15] in conjunction with Remark 2.1 therein, or by adapting [42, Theorem 2.1.14] to the present infinite-dimensional setting, one can deduce that, if $F : \mathcal{H}^2(\mathbb{R}^k) \rightarrow \mathbb{R}$ is convex and its derivative $\nabla F : \mathcal{H}^2(\mathbb{R}^k) \rightarrow \mathcal{H}^2(\mathbb{R}^k)$ in (2.2) is Lipschitz continuous, then for sufficiently small stepsize, the corresponding costs $(J(\alpha^m))_{m \in \mathbb{N}}$ of the sequence $(\alpha^m)_{m \in \mathbb{N}}$ converges to the optimal cost J^* in (1.2) (with $\ell \equiv 0$) with the rate $\mathcal{O}(m^{-1})$. Such a regularity condition holds in particular for commonly used linear-convex MFC problems (see e.g., [12, 1, 31, 47]). For nonconvex F , it is well-known that gradient-based algorithms in general can only find critical points.

However, it is difficult to implement the iteration (2.5) in practice. For each m , since α^m is a non-Markovian stochastic process, the BSDE (2.4) for the adjoint processes (Y^m, Z^m) is typically non-Markovian. This prevents us from evaluating $(\nabla F)(\alpha^m)$ and implementing the updates (2.5) in a pairwise sense as for gradient-based algorithms for deterministic optimal control problems (see e.g., [8]). To be

more precise, let us initialize the iterates (2.5) with $\alpha^0 = 0$. Then one can express (Y^1, Z^1) as $(Y_t^0, Z_t^0) = (u(t, X_t^0), v(t, X_t^0))$ for some deterministic functions u, v (often called decoupling fields), and obtain the stochastic processes (Y^0, Z^0) by computing the functions u, v . Hence, one easily sees that α^1 is a function of X^0 , and consequently, the coefficients of the state dynamics (1.1) and adjoint equations (2.4) for (X^1, Y^1, Z^1) would depend on both X^0 (through α^1) and X^1 . Repeating this process, one observes that (Y^m, Z^m) are functions of time and the enlarged system (X^0, \dots, X^m) , whose computational complexity increases rapidly as the number of iterations grows. A similar difficulty has been observed in [16] for implementing (non-Markovian) Picard iterations to solve coupled FBSDEs.

2.2 Second motivating algorithm: gradient descent in feedback control

In this section, we heuristically derive a gradient descent method for (2.1) by restricting the optimisation over sufficiently regular feedback maps.

To this end, for each sufficiently regular feedback control $\phi : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^k$, let X^ϕ be the corresponding state process satisfying the dynamics: for all $t \in [0, T]$,

$$dX_t = b(t, X_t, \phi(t, X_t), \mathcal{L}_{(X_t, \phi(t, X_t))}) dt + \sigma(t, X_t, \phi(t, X_t), \mathcal{L}_{(X_t, \phi(t, X_t))}) dW_t, \quad X_0 = \xi. \quad (2.6)$$

Then (2.1) (i.e., (1.1)-(1.2) with $\ell \equiv 0$) can be reformulated as minimizing

$$\tilde{J}(\phi) := \int_0^T \int_{\mathbb{R}^d} f(t, y, \phi(t, y), \mathcal{L}_{(X_t^\phi, \phi(t, X_t^\phi))}) \mathcal{L}_{X_t^\phi}(dy) dt + \int_{\mathbb{R}^d} g(y, \mathcal{L}_{X_T^\phi}) \mathcal{L}_{X_T^\phi}(dy) \quad (2.7)$$

over all sufficiently regular feedback controls ϕ . Recall that for any given feedback map ψ , the derivative of \tilde{J} at ϕ in the direction ψ (see e.g., [20, Chapter 6.2.4]) is

$$\left. \frac{d\tilde{J}(\phi + \epsilon\psi)}{d\epsilon} \right|_{\epsilon=0} = \int_0^T \int_{\mathbb{R}^d} \langle \delta \tilde{J}_\phi(t, y), \psi(t, y) \rangle \mathcal{L}_{X_t^\phi}(dy) dt, \quad (2.8)$$

where $\delta \tilde{J}_\phi : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^k$ is given by

$$\begin{aligned} \delta \tilde{J}_\phi(t, x) &= (\partial_a \mathbb{H})(t, x, \phi(t, x), \mathcal{L}_{(X_t^\phi, \phi(t, X_t^\phi))}, (\nabla_x u^\phi)(t, x), (\text{Hess}_x u^\phi)(t, x)) \\ &\quad + \mathbb{E} \left[(\partial_\nu \mathbb{H})(t, \tilde{X}_t^\phi, \phi(t, \tilde{X}_t^\phi), \mathcal{L}_{(X_t^\phi, \phi(t, X_t^\phi))}, (\nabla_x u^\phi)(t, \tilde{X}_t^\phi), (\text{Hess}_x u^\phi)(t, \tilde{X}_t^\phi)) (x, \phi(t, x)) \right], \end{aligned} \quad (2.9)$$

\tilde{X}^ϕ is an independent copy of \tilde{X}^ϕ , $\mathbb{H} : [0, T] \times \mathbb{R}^d \times \mathbb{R}^k \times \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^k) \times \mathbb{R}^d \times \mathbb{R}^{d \times d} \rightarrow \mathbb{R}$ is given by

$$\mathbb{H}(t, x, a, \eta, y, z) := \langle b(t, x, a, \eta), y \rangle + \frac{1}{2} \text{tr}((\sigma \sigma^\top)(t, x, a, \eta) z) + f(t, x, a, \eta),$$

and $u^\phi : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$ satisfies the following linear PDE:

$$\begin{aligned} (\partial_t u)(t, x) &= -\frac{\delta}{\delta \mu} \left[\int_{\mathbb{R}^d} \mathbb{H}(t, y, \phi(t, y), \mathcal{L}_{(X_t^\phi, \phi(t, X_t^\phi))}, (\nabla_x u)(t, y), (\text{Hess}_x u)(t, y)) \mathcal{L}_{X_t^\phi}(dy) \right] (x), \\ u(T, x) &= \frac{\delta}{\delta \mu} \left[\int_{\mathbb{R}^d} g(y, \mathcal{L}_{X_T^\phi}) \mathcal{L}_{X_T^\phi}(dy) \right] (x), \end{aligned} \quad (2.10)$$

with $\frac{\delta}{\delta \mu}$ being the linear functional derivative with respect to the state law $(\mathcal{L}_{X_t^\phi})_{t \in [0, T]}$ (see [20, Definition 5.43]).

Now gradient descent updates of feedback controls can be derived by selecting a function representation of the directional derivative given in (2.8). For instance, let ϕ^0 be an initial feedback control, and $\tau > 0$ be a stepsize. For all $m \in \mathbb{N} \cup \{0\}$, suppose that the law $\mathcal{L}_{X_t^{\phi^m}}$ admits a density $\frac{d\mathcal{L}_{X_t^{\phi^m}}}{dx}$ with respect to the Lebesgue measure for almost everywhere $t \in (0, T]$, then (2.8) can be written as:

$$\left. \frac{d\tilde{J}(\phi^m + \epsilon\psi)}{d\epsilon} \right|_{\epsilon=0} = \int_0^T \int_{\mathbb{R}^d} \left\langle \delta \tilde{J}_{\phi^m}(t, x), \frac{d\mathcal{L}_{X_t^{\phi^m}}}{dx}(x), \psi(t, x) \right\rangle dx dt,$$

which allows for interpreting the function $(t, x) \mapsto \delta \tilde{J}_{\phi^m}(t, x) \frac{d\mathcal{L}_{X_t^{\phi^m}}}{dx}(x)$ as the gradient of \tilde{J} at ϕ^m with respect to the canonical inner product on $L^2([0, T] \times \mathbb{R}^d)$. This suggests updating ϕ^m by

$$\phi^{m+1}(t, x) = \phi^m(t, x) - \tau \delta \tilde{J}_{\phi^m}(t, x) \frac{d\mathcal{L}_{X_t^{\phi^m}}}{dx}(x), \quad (t, x) \in [0, T] \times \mathbb{R}^d, \quad (2.11)$$

which is a direct application of gradient descent for (2.1) (or equivalently (1.1)-(1.2) with $\ell \equiv 0$) over feedback controls.

Pros and cons of Iteration (2.11). Iteration (2.11) directly updates the feedback controls, and hence its complexity at each iteration remains the same as the iteration proceeds. It overcomes the drawback of Iteration (2.5), whose computational complexity increases as the number of iterations grows as illustrated at the end of Section 2.1.

However, Iteration (2.11) has several critical deficiencies. Compared with Iteration (2.5), the gradient direction in (2.11) involves the density of the state process, which may not exist and can be computationally expensive to evaluate if the density function is irregular. This irregularity arises in particular when the diffusion coefficient of (2.6) degenerates or when the initial state ξ has a singular density. Moreover, extending Iteration (2.11) to general cost functionals (1.2) with nonzero ℓ is more challenging. Observe that in the open-loop formulation, $\alpha \mapsto \mathbb{E}[\int_0^T \ell(\alpha_t) dt]$ in (1.2) is convex due to the convexity of ℓ and hence can be easily handled by the associated proximal map; see Section 2.4 for more details. However, due to the nonlinear dependence of X^ϕ on ϕ , $\phi \mapsto \mathbb{E}[\int_0^T \ell(\phi(t, X_t^\phi)) dt]$ is typically nonconvex and evaluating the associated proximal map is computationally expensive.

Furthermore, analysing the convergence of (2.11) is technically more challenging. This is because the mapping $\phi \mapsto \tilde{J}(\phi)$ in (2.7) is typically nonconvex with respect to feedback controls, even for deterministic linear-quadratic problems without mean-field interaction (see [30, Proposition 2.4] for a concrete example). This implies that the convergence analysis of such PGMs is linked to analyzing nonasymptotic performance of gradient search for nonconvex objectives, which has always been one of the formidable challenges in optimization theory.

In fact, it is unclear how to choose a suitable function space to analyze Iteration (2.11). By (2.9) and (2.11), the regularity of ϕ^{m+1} depends on the regularity of $\text{Hess}_x u^{\phi^m}$. As u^{ϕ^m} satisfies (2.10) whose coefficients involve ϕ^m , it is expected that the regularity of $\text{Hess}_x u^{\phi^m}$ is controlled by the regularity of $\text{Hess}_x \phi^m$. This suggests that estimating the derivatives of ϕ^{m+1} requires bounds on higher order derivatives of ϕ^m , and it is unclear how to close this norm gap.

2.3 A heuristic modification of gradient methods in feedback control

This section proposes a novel iterative method for MFC problems with regular running costs ($\ell \equiv 0$ in (1.2)) by combining the features of Iterations (2.5) and (2.11). On one hand, the algorithm directly updates the feedback controls as Iteration (2.11), and hence the computational complexity of the adjoint variables do not change with respect to the number of iterations. On the other hand, it heuristically replaces the gradient direction over feedback controls by the functional derivative in Iteration (2.5). In particular, at the m -th iteration, we will update the present feedback map by evaluating the functional derivative of $J(\cdot; \xi_0)$ at the open-loop control α^{ϕ^m} induced by the current policy ϕ^m , and then obtain the update direction based on a Markovian representation of the gradient. The latter can be seen as a simplified version of Iteration (2.11) that does not require computing the state process's density, the derivative of feedback functions, or the prox operator in the space of feedback functions in the case of non-smooth costs (see Remarks 2.1 and 2.2, respectively).

We shall represent the control α^m as $\alpha_t^m = \phi^m(t, X_t^m)$ with some deterministic feedback function $\phi^m : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^k$ at each iteration. The update direction of ϕ^m is obtained via the functional derivative in (2.5). More precisely, at the m -th iteration, given ϕ^m , consider the associated the controlled state dynamics: $X_0 = \xi$, and for all $t \in [0, T]$,

$$dX_t = b(t, X_t, \phi^m(t, X_t), \mathcal{L}_{(X_t, \phi^m(t, X_t))}) dt + \sigma(t, X_t, \phi^m(t, X_t), \mathcal{L}_{(X_t, \phi^m(t, X_t))}) dW_t. \quad (2.12)$$

Let X^m be the solution of (2.12); we then seek the adjoint processes (Y^m, Z^m) satisfying (2.4) with (X^α, α) replaced by $(X^m, \phi^m(\cdot, X^m))$, in order to evaluate the gradient of F at the corresponding control

process $\phi^m(\cdot, X^m)$ as in (2.2). The feedback structure of α^m implies that there exist deterministic decoupling fields $u^m : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $v^m : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^{d \times n}$ such that

$$Y_t^m = u^m(t, X_t^m), \quad Z_t^m = v^m(t, X_t^m), \quad dt \otimes d\mathbb{P}\text{-a.e.} \quad (2.13)$$

Moreover, the functional derivative $\nabla F(\alpha^m)$ admits a Markovian representation $(\nabla F)(\phi^m(\cdot, X^m))_t = (\nabla F)(\phi^m)(t, X_t^m)$, $dt \otimes d\mathbb{P}\text{-a.e.}$, with the function $(\nabla F)(\phi^m) : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^k$ defined by

$$\begin{aligned} & (\nabla F)(\phi^m)(t, x) \\ &:= (\partial_a H)(t, x, \phi^m(t, x), \mathcal{L}_{(X_t^m, \phi^m(t, X_t^m))}, u^m(t, x), v^m(t, x)) \\ &+ \tilde{\mathbb{E}}[(\partial_\nu H)(t, \tilde{X}_t^m, \phi^m(t, \tilde{X}_t^m), \mathcal{L}_{(X_t^m, \phi^m(t, X_t^m))}, u^m(t, \tilde{X}_t^m), v^m(t, \tilde{X}_t^m))(x, \phi^m(t, x))], \end{aligned} \quad (2.14)$$

where H is defined in (2.3). Hence, once the functions (u^m, v^m) are determined, one can compute the Markovian representation (2.14) of the functional derivative $(\nabla F)(\phi^m)$, along with a gradient descent step can be applied to obtain the updated feedback map ϕ^{m+1} ; see (2.20).

To compute the decoupling fields (u^m, v^m) in (2.13) for each m , we connect them with solutions to some PDE system. In fact, the nonlinear Feynman–Kac formula in [44] shows that, if u^m is sufficiently smooth, then $v^m(t, x) = (\nabla_x u^m)(t, x) \sigma(t, x, \mathcal{L}_{X_t^m})$ for all $(t, x) \in [0, T] \times \mathbb{R}^d$, and u^m solves the following system of parabolic linear PDEs (depending on ϕ^m and the law of X^m):

$$(\partial_t u)(t, x) + (\mathcal{L}^m u)(t, x) = -f^m(t, x, u, v), \quad (2.15)$$

with $u(T, x) = h^m(x)$ and $v(t, x) = (\nabla_x u)(t, x) \sigma^m(t, x)$ for all $(t, x) \in [0, T] \times \mathbb{R}^d$, where \mathcal{L}^m is the (vector-valued) differential operator such that for each $\varphi = (\varphi_1, \dots, \varphi_d) \in C^{1,2}([0, T] \times \mathbb{R}^d; \mathbb{R}^d)$, $i \in \{1, \dots, d\}$ and $(t, x) \in [0, T] \times \mathbb{R}^d$,

$$(\mathcal{L}^m \varphi)_i(t, x) = \frac{1}{2} \text{tr}(\sigma^m(t, x) \sigma^m(t, x)^\top (\text{Hess}_x \varphi_i)(t, x)) + \langle b^m(t, x), (\nabla_x \varphi_i)(t, x) \rangle, \quad (2.16)$$

with $b^m : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $\sigma^m : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^{d \times n}$ given by

$$\begin{aligned} b^m(t, x) &= b(t, x, \phi^m(t, x), \mathcal{L}_{(X_t^m, \phi^m(t, X_t^m))}), \\ \sigma^m(t, x) &= \sigma(t, x, \phi^m(t, x), \mathcal{L}_{(X_t^m, \phi^m(t, X_t^m))}), \end{aligned} \quad (2.17)$$

and the functions f^m and h^m satisfy for all $(t, x) \in [0, T] \times \mathbb{R}^d$ (cf. (2.4)),

$$\begin{aligned} f^m(t, x, u, v) &= (\partial_x H)(t, x, \phi^m(t, x), \mathcal{L}_{(X_t^m, \phi^m(t, X_t^m))}, u(t, x), v(t, x)) \\ &+ \tilde{\mathbb{E}}[(\partial_\mu H)(t, \tilde{X}_t^m, \phi^m(t, \tilde{X}_t^m), \mathcal{L}_{(X_t^m, \phi^m(t, X_t^m))}, u(t, \tilde{X}_t^m), v(t, \tilde{X}_t^m))(x, \phi^m(t, x))], \end{aligned} \quad (2.18)$$

$$h^m(x) = (\partial_x g)(x, \mathcal{L}_{X_T^m}) + \tilde{\mathbb{E}}[(\partial_\mu g)(\tilde{X}_T^m, \mathcal{L}_{X_T^m})(x)]. \quad (2.19)$$

Algorithm 1 summarizes the above gradient descent method for (1.1)–(1.2) (with $\ell \equiv 0$), which will be referred to as the iterative PDE-based (IPDE) method hereafter.

Algorithm 1 Iterative PDE-based method for MFC problems

- 1: **Input:** Choose the initial feedback control $\phi^0 : [0, T] \times \mathbb{R}^d \mapsto \mathbb{R}^k$ and stepsize $\tau > 0$.
- 2: **for** $m = 0, 1 \dots$ **do**
- 3: Compute the law of the state process X^m governed by (2.12).
- 4: Compute the decoupling fields (u^m, v^m) of (Y^m, Z^m) by solving (2.15) based on the law of X^m .
- 5: Update the controls such that for all $(t, x) \in [0, T] \times \mathbb{R}^d$,

$$\phi^{m+1}(t, x) = \phi^m(t, x) - \tau (\nabla F)(\phi^m)(t, x) \quad (2.20)$$

with $(\nabla F)(\phi^m)$ defined by (2.14).

- 6: **end for**
-

Remark 2.1. For the avoidance of doubt, the update (2.20) is neither a gradient step for the stochastic formulation of the problem using a Markovian representation of the open-loop controls, nor is it a gradient

step with respect to the feedback map. It is not the former because the update (2.20) is not equivalent to the corresponding gradient step (2.5), since subtracting feedback controls is not equivalent to subtracting the corresponding stochastic processes. It is not the latter because the update direction does not involve the density of $X_t^{\phi^m}$, and (u^m, v^m) in (2.14) is different from $(\nabla_x u^{\phi^m}, \text{Hess}_x u^{\phi^m})$ in (2.9), as they are solutions to different PDEs. To see this, assume for simplicity that all coefficients are independent of the measure components. Then the PDE (2.15) for u^m simplifies to

$$(\partial_t u)(t, x) + (\mathcal{L}^m u)(t, x) = -(\partial_x H)(t, x, \phi^m(t, x), u(t, x), (\nabla_x u)(t, x) \sigma^m(t, x)). \quad (2.21)$$

On the other hand, formally differentiating the PDE (2.10) for u^{ϕ^m} with respect to x suggests that $\nabla_x u^{\phi^m}$ satisfies the following equation:

$$(\partial_t u)(t, x) + (\mathcal{L}^m u)(t, x) = -(\partial_x \mathbb{H}^m)(t, x, u(t, x), (\nabla_x u)(t, x)), \quad (2.22)$$

where $\mathbb{H}^m : [0, T] \times \mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ is defined by

$$\mathbb{H}^m(t, x, y, z) := \langle b(t, x, \phi^m(t, x)), y \rangle + \frac{1}{2} \text{tr}((\sigma \sigma^\top)(t, x, \phi^m(t, x)) z) + f(t, x, \phi^m(t, x)).$$

It is clear from the chain rule that $\partial_x \mathbb{H}^m$ consists of both the term $\partial_x H$ in (2.21) and an additional term involving the derivative $\partial_x \phi^m$.

The above observation indicates that the modified gradient update (2.20) imposes weaker regularity requirements on the iterate ϕ^m than the vanilla gradient descent (2.11), as it avoids evaluating the density of X^{ϕ^m} and the derivatives of ϕ^m in the computation of the gradient $(\nabla F)(\phi^m)$. This difference in gradient evaluation not only enhances the stability of (2.20) (by preventing numerical instability when computing an irregular density and differentiating the feedback controls), but is also essential for analyzing the convergence of (2.20); see Section 2.5 for details.

Concrete numerical methods for solving (2.15) (and consequently the functions (u^m, v^m)) will be given in Sections 3 and 4. Here, we point out the following three features of the PDE system (2.15), which are crucial in the design of numerical methods: (i) As the Hamiltonian H defined in (2.3) is affine in the components y and z , the function f^m is linear in u and v , and hence (2.15) is a linear PDE system. (ii) Due to the measure dependence in f^m , (2.15) is nonlocal in the sense that the value of the solution u^m at each point evolves based on the weighted average of other values of u^m with respect to the marginal laws $(\mathcal{L}_{X_t^m})_{t \in [0, T]}$ of the process X^m . (iii) Even though the i -th component of the differential operator \mathcal{L}^m in (2.16) only involves the i -th component of the solution u^m , the function f^m in (2.15) in general results in a coupling among all components of u and their gradients v . In the special cases where the diffusion coefficient σ of (1.1) depends only on time (see Section 5 for concrete examples), the function f^m is independent of v and hence the system (2.15) is only coupled through the solution u^m .

We emphasize that, a key feature of the IPDE method is that it computes the functions (u^m, v^m) for each $(t, x) \in [0, T] \times \mathbb{R}^d$ based on a PDE formulation, in contrast to the pure data-driven algorithms in [6, 21, 28, 48], which solve (u^m, v^m) merely along the trajectories of X^m . Hence, the IPDE method leads to a more accurate approximation of the optimal feedback control, especially outside the support of the optimal state process of (1.1)-(1.2). In particular, the IPDE method is capable of recovering important structural features of the optimal control, and the performance of the approximate feedback control is robust with respect to perturbation of model parameters; see Section 5 for a detailed comparison between the (F)IPDE method and several data-driven algorithms.

2.4 A heuristic acceleration of gradient methods with nonsmooth cost

This section extends Algorithm 1 in Section 2.3 to MFC problems (1.1)-(1.2) with nonsmooth costs and further accelerates the algorithm convergence by incorporating an additional moment step. To handle the nonsmooth term $\ell : \mathbb{R}^k \rightarrow \mathbb{R} \cup \{\infty\}$, for each $\tau > 0$, we consider the following proximal map $\text{prox}_{\tau \ell} : \mathbb{R}^k \rightarrow \mathbb{R}^k$:

$$\text{prox}_{\tau \ell}(a) = \arg \min_{z \in \mathbb{R}^k} \left(\frac{1}{2} |z - a|^2 + \tau \ell(z) \right), \quad a \in \mathbb{R}^k,$$

which is well-defined as ℓ is proper, lower semicontinuous and convex. For many practically important nonsmooth functions ℓ , the proximal function $\text{prox}_{\tau \ell}$ can be evaluated either analytically (see e.g., [13,

Chapter 6]) or approximately by efficient numerical methods (see [51] and the references therein). The acceleration of (2.20) is inspired by Nesterov's accelerated proximal gradient (NAG) method (also known as the Fast Iterative Shrinkage-Thresholding Algorithm [15]) tailored to the present setting. The algorithm is summarized as follows and will be referred to as the fast iterative PDE-based (FIPDE) method.

Algorithm 2 Fast Iterative PDE-based method for MFC problems

- 1: **Input:** Choose the initial feedback control $\phi^0 : [0, T] \times \mathbb{R}^d \mapsto \mathbb{R}^k$ and stepsize $\tau > 0$. Set $\psi^0 = \phi^0$.
- 2: **for** $m = 0, 1 \dots$ **do**
- 3: Compute the law of the state process X^m governed by (2.12).
- 4: Compute the decoupling fields (u^m, v^m) of (Y^m, Z^m) by solving (2.15) based on the law of X^m .
- 5: Update the controls such that for all $(t, x) \in [0, T] \times \mathbb{R}^d$,

$$\phi^{m+1}(t, x) = \text{prox}_{\tau\ell}(\psi^m(t, x) - \tau(\nabla F)(\psi^m)(t, x)), \quad (2.23a)$$

$$\psi^{m+1}(t, x) = \phi^{m+1}(t, x) + \frac{m}{m+3}(\phi^{m+1}(t, x) - \phi^m(t, x)), \quad (2.23b)$$

with $(\nabla F)(\psi^m)$ defined by (2.14).

- 6: **end for**
-

The ratio $\frac{m}{m+3}$ in (2.23b) is often referred to as the momentum parameter, whose selection adheres to the general guideline provided by [15].³

Remark 2.2. Further to Remark 2.1 for Algorithm 1, Algorithm 2 is a heuristic combination of the NAG method applied to the stochastic formulation of the control problem over open-loop controls, and to the deterministic reformulation over feedback controls discussed in Section 2.2. The computation of the prox in (2.23a) is the one for the formulation of the problem with stochastic processes. It aims to overcome the difficulty in computing a suitable prox for feedback maps caused by the lack of convexity of $\phi \mapsto \mathbb{E}[\int_0^T \ell(\phi(t, X_t^\phi)) dt]$ (see the end of Section 2.2). In particular, we have exploited the structure of the nonsmooth functional $G : \alpha \mapsto \mathbb{E}[\int_0^T \ell(\alpha_t) dt]$ and explicitly expressed the proximal of $G : \mathcal{H}^2(\mathbb{R}^k) \rightarrow \mathbb{R} \cup \{\infty\}$ via a pointwise composition of the proximal of ℓ . The addition/subtraction of controls in (2.23a) and (2.23b) is the addition/subtraction of feedback maps and is not equivalent to the addition/subtraction of the corresponding processes.

In practice, we represent the feedback maps $(\phi^m, \psi^m)_{m \in \mathbb{N}}$ and the decoupling fields $(u^m, v^m)_{m \in \mathbb{N}}$ in suitable parametric forms, whose precise choices depend on the dimension of the problem and the numerical methods used to solve the PDE system (2.15) (see Sections 3 and 4 for details). Given a parameterized feedback map ψ^m , the law of the controlled state process X^m in Step 3 of Algorithm 2 can be approximated by a particle method and an Euler–Maruyama discretization of (2.12). For instance, let $N \in \mathbb{N}$ be the number of particles and $\{0 = t_0 < \dots < t_M = T\}$ a partition of $[0, T]$ with time stepsize $\Delta t = T/M$ for some $M \in \mathbb{N}$. Then we consider the discrete-time interacting particle system $(X_t^{l,N})_{t \in [0, T]}$, $l = 1, \dots, N$, such that $X_0^{l,N} = \xi^l$, for $j = 0, \dots, M-1$ and $t \in [t_j, t_{j+1})$, $X_t^{l,N} = X_{t_j}^{l,N}$, and

$$X_{t_{j+1}}^{l,N} = X_{t_j}^{l,N} + b(t_j, X_{t_j}^{l,N}, \psi^m(t_j, X_{t_j}^{l,N}), \mu_{N,t_j}^{X, \psi^m})\Delta t + \sigma(t_j, X_{t_j}^{l,N}, \psi^m(t_j, X_{t_j}^{l,N}), \mu_{N,t_j}^{X, \psi^m})\Delta W_j^l, \quad (2.24)$$

with $\mu_{N,t_j}^{X, \psi^m}(dx) = \frac{1}{N} \sum_{l=1}^N \delta_{(X_{t_j}^{l,N}, \psi^m(t_j, X_{t_j}^{l,N}))}(dx)$, where $(\xi^l)_{l=1, \dots, N}$ and $(W^l)_{l=1, \dots, N}$ are independent copies of ξ and W , respectively, $\Delta W_j^l = W_{t_{j+1}}^l - W_{t_j}^l$ for all j, l , and δ_x denotes the Dirac measure supported at x for all $x \in \mathbb{R}^d$.

Compared with the deterministic methods in [2, 3, 4, 22], (2.24) avoids the numerical challenge in solving the nonlinear FP equation, and allows for efficient computation of the marginal distribution of the state process, especially for MFC problems (1.1)–(1.2) with degenerate diffusion coefficient σ or initial state ξ with singular density. In fact, it is well-known that, for sufficiently regular feedback maps ψ^m , the

³ By [15, Theorem 4.4], NAG method (or FISTA) with the momentum step $y^{m+1} = x^{m+1} + \frac{\theta_{m+1}-1}{\theta_{m+2}}(x^{m+1} - x^m)$ converges for convex optimisation problems, provided that $\theta_{m+1} \geq 1$ and $\theta_m^2 \geq \theta_{m+1}^2 - \theta_{m+1}$ for all $m \geq 0$. For all $m \in \mathbb{N}$, (2.23b) corresponds to $\theta_m = (m+1)/2$, while [15] chooses $\theta_{m+1} \geq 1$ such that $\theta_m^2 = \theta_{m+1}^2 - \theta_{m+1}$.

empirical measure μ_N^{X, ψ^m} converges to the law of $(X^m, \psi^m(\cdot, X^m))$ in the Wasserstein metric as M, N tend to infinity; see also [18, 10] for the convergence rates of (2.24) in terms of M and N . For our numerical experiments in Section 5, we choose sufficiently large M and N such that the presented results are not influenced by those choices.

As numerical approximations of (2.12) are relatively well-understood, in the subsequent sections, we focus on the numerical approximations of the PDE system (2.15) with given parameterized feedback function ψ^m and empirical approximations $\mu_{N, \cdot}^{X, \psi^m}$ of the law of $(X^m, \psi^m(\cdot, X^m))$. In particular, we shall propose a class of monotone schemes in Section 3 for (2.15) with spatial dimension $d \leq 3$, and neural network-based schemes in Section 4 for the high-dimensional setting.

2.5 Convergence of Algorithms 1 and 2

We begin by recalling the discussion of convergence for the NAG method with open-loop controls at the end of Section 2.1. Furthermore, as shown in Theorem 4.4 of [15], including a momentum step after (2.5) yields an improved convergence rate $\mathcal{O}(m^{-2})$ for minimising smooth convex functionals J , which is the optimal convergence rate of gradient-based algorithms (see [15] and references therein). The same convergence rate also holds when convex nonsmooth costs are incorporated through an application of a proximal operator. Additionally, if the controls are updated with approximate proximal operators, one can still recover the same convergence rate, provided that the errors made in the calculation of the proximal operator decrease at appropriate rates [51]. For nonconvex functionals, it is known that the momentum step can produce highly oscillatory solutions. Additional monitoring steps are typically necessary to ensure a sufficient descent of the function values and the convergence of the iterates to critical points [14, 39, 29, 40].

Unfortunately, existing convergence results of the NAG method cannot be applied to Algorithm 2. As pointed out in Remarks 2.1 and 2.2, Algorithm 2 is not a direct application of the NAG method (or FISTA) to the stochastic formulation of the problem (1.1)-(1.2) nor to the deterministic reformulation discussed in Section 2.2. The gradient direction and the proximal map in (2.23a) are computed using the stochastic formulation over open-loop controls, whereas the updated policy is obtained forming an addition of feedback functions which, in general, is different to the corresponding combination of open-loop controls. Hence, although our numerical results in Section 5.2 indicate that Algorithm 2 performs reasonably well even for nonconvex control problems, a convergence analysis of Algorithm 2 remains an open question.

Note that a comparable discrepancy between two formulations arises even in the absence of the momentum step (2.23b) (i.e., $\psi^m = \phi^m$ for all m). In the recent work [49], such an algorithm has been analyzed via a contraction argument for specific drift-controlled nonsmooth control problems without mean-field interaction. We identify conditions under which the algorithm generates uniform Lipschitz feedback controls $(\phi^m)_{m \in \mathbb{N}}$, and further establish that the associated control processes converge to a critical point of the functional J in (1.2). The convergence result holds when the diffusion coefficient is uncontrolled, and in addition one of the following five cases is satisfied: (i) time horizon T is small; (ii) running cost is sufficiently convex in control; (iii) costs depend weakly on state; (iv) control affects state dynamics weakly; (v) state dynamics is strongly dissipative. Note that the conditions allow for nonlinear state dynamics with degenerate noise, and nonconvex and nonsmooth cost functions. We refer the reader to [49] for rigorous statements of the convergence results.

The main result in [49] shows that in the case without mean-field interaction and without the momentum step, convergence of Algorithm 2 to stationary points can be proven, despite the aforementioned mismatch appearing in our formulation. This in particular shows the convergence of Algorithm 1 without mean-field interaction. The proof exploits regularity estimates for the decoupling fields $(u^m, v^m)_{m \in \mathbb{N}}$ in (2.13) and does not require the convexity of the cost landscape. We conjecture that similar convergence results also hold in the mean-field setting, whose rigorous proof require establishing the regularity of the decoupling fields in the measure component, and is left for future research. However, at the present stage, it is unclear if this mismatch when applying an additional momentum step (2.23b) in the policy update can be overcome, even under the assumptions of [49].

3 Implementation of the FIPDE method via finite differences

In this section, we discuss the practical implementation of the FIPDE method in Section 2 for low-dimensional state dynamics (1.1). In particular, we shall propose a class of semi-implicit monotone finite difference approximations for solving (2.15) at the m -th NAG iteration, which achieves an efficient performance in terms of the computation time if (1.1) has a spatial dimension $d \leq 2$.

Throughout this section, we focus on the m -th NAG iteration with $m \in \mathbb{N} \cup \{0\}$ and assume the feedback function ψ^m and particle approximation $(X^{l,N})_{l=1,\dots,N}$ of X^m are given (see the discussion below Algorithm 2). Then, we need to solve the following nonlocal parabolic PDE system (cf. (2.15)):

$$(\partial_t u)(t, x) + (\mathcal{L}_N^m u)(t, x) = -f_N^m(t, x, u, v), \quad (3.1)$$

with $u(T, x) = h^m(x)$ and $v(t, x) = (\nabla_x u)(t, x)\sigma^m(t, x)$ for all $(t, x) \in [0, T] \times \mathbb{R}^d$, where the operator \mathcal{L}_N^m (resp. the function f_N^m) is defined similar to (2.16) (resp. (2.18)), but depends on the empirical measure $\mu_{N,t}^{X,\psi^m}(dx) = \frac{1}{N} \sum_{l=1}^N \delta_{(X_t^{l,N}, \psi^m(t, X_t^{l,N}))}(dx)$ instead of the law $\mathcal{L}_{(X_t^m, \psi^m(t, X_t^m))}$ for all $t \in [0, T]$. To simplify the presentation, we shall focus on the uniform spatial grid $\{x_k\}_k = h\mathbb{Z}^d$ on \mathbb{R}^d with mesh size $h > 0$ and a time partition $\{t_j\}_{j=0}^M$ with time stepsize $\Delta t = T/M$ for $M \in \mathbb{N}$, but similar schemes can be designed for unstructured nondegenerate grids as well.

We start by introducing a semi-implicit timestepping approximation to (3.1). Observe that the i -th component of the differential operator \mathcal{L}_N^m depends only on the i -th component of the solution u , and all nonlocal and coupling terms appear in f_N^m . Hence, we shall adopt implicit timestepping for the local operator \mathcal{L}_N^m and explicit timestepping for f_N^m , which leads to the following (backward) time discretization of (3.1): $U^M(x) = h^m(x)$ for all $x \in \mathbb{R}^d$ and for all $j = 1, \dots, M$, $x \in \mathbb{R}^d$,

$$\frac{U^j(x) - U^{j-1}(x)}{\Delta t} + (\mathcal{L}_N^m U^{j-1})(x) = -f_N^m(t_j, x, U^j, V^j),$$

where for all $x \in \mathbb{R}^d$, $U^j(x)$ is the approximation of $u(t_j, x)$ and $V^j(x) = (\nabla_x U^j)(x)\sigma_N^m(t_j, x)$. Note that the implicit timestepping for \mathcal{L}_N^m enables us to enjoy a less restrictive stability condition than that for fully explicit schemes, while the explicit timestepping for f_N^m avoids solving the dense system resulting from the mean-field terms, and allows us to solve for each component of U^j independently, given U^{j+1} .

We proceed to perform spatial discretization of \mathcal{L}_N^m . Note that the i -th component of \mathcal{L}_N^m depends only on the i -th component of u , and all components of \mathcal{L}_N^m have the same coefficients (see (2.16) and (2.17)). Then, as shown in [17], one can construct monotone and consistent approximations of \mathcal{L}_N^m such that for any $\varphi = (\varphi_1, \dots, \varphi_d) \in C^2(\mathbb{R}^d, \mathbb{R}^d)$, $i = 1, \dots, d$, $j = 0, \dots, M$, $k \in \mathbb{Z}^d$, $L_{N,h}^{m,j}[\varphi]_k = (L_{N,h}^{m,j}[\varphi]_{1,k}, \dots, L_{N,h}^{m,j}[\varphi]_{d,k})$ satisfies

$$\begin{aligned} L_{N,h}^{m,j}[\varphi]_{i,k} &= \sum_{q \in \mathbb{Z}^d} a_{h,j,q,k}^{m,N} [\varphi_i(x_q) - \varphi_i(x_k)], \\ |(\mathcal{L}_N^m \varphi)_i(t_j, x_k) - L_{N,h}^{m,j}[\varphi]_{i,k}| &\rightarrow 0, \quad \text{as } h \rightarrow 0, \end{aligned} \quad (3.2)$$

with coefficients $a_{h,j,q,k}^{m,N} \geq 0$ for all q, k, j . The precise construction of such numerical approximations depends on the structures of the coefficients b^m and σ^m . In particular, one can adopt the standard finite difference schemes in [11, 37] if the diffusion coefficient $\sigma^m(\sigma^m)^\top$ is diagonally dominant, and use the semi-Lagrangian scheme in [25] for general cases. We refer the reader to Section 5 for more details.

It remains to discretize the term $f_N^m(t, x, u, v)$. By (2.3) and (2.18), we have

$$f_N^m(t, x, u, v) = f_N^{m,\text{re}}(t, x, u) + f_N^{m,\text{ex}}(t, x, u),$$

with the terms

$$\begin{aligned} f_N^{m,\text{re}}(t, x, u) &:= ((\partial_x b)(t, x, \psi^m(t, x), \mu_{N,t}^{X,\psi^m}))^\top u(t, x) + (\partial_x f)(t, x, \psi^m(t, x), \mu_{N,t}^{X,\psi^m}) \\ &\quad + \mathbb{E}^{\mu_{N,t}^X} \left[((\partial_\mu b)(t, \cdot, \psi^m(t, \cdot), \mu_{N,t}^{X,\psi^m})(t, \psi^m(t, x)))^\top u(t, \cdot) \right. \\ &\quad \left. + (\partial_\mu f)(t, \cdot, \psi^m(t, \cdot), \mu_{N,t}^{X,\psi^m})(t, \psi^m(t, x)) \right], \\ f_N^{m,\text{ex}}(t, x, u) &:= \left((\nabla_x \sigma)(t, x, \psi^m(t, x), \mu_{N,t}^{X,\psi^m}) \right)^\top \left((\nabla_x u)(t, x) \sigma(t, x, \psi^m(t, x), \mu_{N,t}^{X,\psi^m}) \right) \end{aligned}$$

$$+ \mathbb{E}^{\mu_{N,t}^X} \left[\left((\partial_\mu \sigma)(t, \cdot, \psi^m(t, \cdot), \mu_{N,t}^{X, \psi^m})(t, \psi^m(t, x)) \right)^\top \left((\nabla_x u)(t, \cdot) \sigma(t, \cdot, \psi^m(t, \cdot), \mu_{N,t}^{X, \psi^m}) \right) \right],$$

where $\mathbb{E}^{\mu_{N,t}^X}[\varphi(\cdot)] := \frac{1}{N} \sum_{l=1}^N \varphi(X_t^{l,N})$ for given function $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}^d$, and $f_N^{m, \text{ex}}$ uses $v(t, x) = (\nabla_x u)(t, x) \sigma^m(t, x)$. Note that if σ is independent of the state variable and the marginal law of the state variable (see Section 5 for concrete examples), then $f_N^{m, \text{ex}}(t, x, u) \equiv 0$ and hence f_N^m is independent of the gradients of u .

Now let $(U_k^j)_{k \in \mathbb{Z}^d}$ be a discrete approximation of $u(t_j, \cdot)$ on the grid $\{x_k\}_{k \in \mathbb{Z}^d}$. We approximate $f_N^{m, \text{re}}$ by replacing $u(t_j, \cdot)$ with the monotone interpolation of $(U_k^j)_{k \in \mathbb{Z}^d}$:

$$\begin{aligned} f_{N,h}^{m, \text{re}}(t_j, x_k, U^j) &:= ((\partial_x b)(t_j, x_k, \psi^m(t_j, x_k), \mu_{N,t_j}^{X, \psi^m}))^\top U_k^j + (\partial_x f)(t_j, x_k, \psi^m(t_j, x_k), \mu_{N,t_j}^{X, \psi^m}) \\ &\quad + \mathbb{E}^{\mu_{N,t_j}^X} \left[((\partial_\mu b)(t_j, \cdot, \psi^m(t_j, \cdot), \mu_{N,t_j}^{X, \psi^m})(t_j, \psi^m(t_j, x_k)))^\top \mathbf{i}_h[U^j](\cdot) \right. \\ &\quad \left. + (\partial_\mu f)(t_j, \cdot, \psi^m(t_j, \cdot), \mu_{N,t_j}^{X, \psi^m})(t_j, \psi^m(t_j, x_k)) \right], \end{aligned} \quad (3.3)$$

where \mathbf{i}_h is the piecewise linear/multilinear interpolation operator such that for all $\varphi = (\varphi_1, \dots, \varphi_d) : h\mathbb{Z}^d \rightarrow \mathbb{R}^d$, $i = 1, \dots, d$, $x \in \mathbb{R}^d$,

$$(\mathbf{i}_h[\varphi])_i(x) = \sum_{k \in \mathbb{Z}^d} \varphi_i(x_k) \omega_k(x; h),$$

with the standard ‘‘tent functions’’ $\{\omega_k\}_k$ satisfying $0 \leq \omega_k(x; h) \leq 1$, $\sum_k \omega_k = 1$, $\omega_k(x_j; h) = \delta_{kj}$ ⁴ and $\text{supp } \omega_k \subset \{x \in \mathbb{R}^d \mid |x - x_k| \leq 2h\}$. To approximate $f_N^{m, \text{ex}}$, we observe that

$$f_N^{m, \text{ex}}(t, x, u) = \sum_{i,l=1}^d \left(b_{N,il}^m(t, x) (\partial_{x_l} u_i)(t, x) + \mathbb{E}^{\mu_{N,t}^X} [c_{N,il}^m(t, x, \cdot) (\partial_{x_l} u_i)(t, \cdot)] \right), \quad (3.4)$$

for some functions $b_{N,il}^m : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$, $c_{N,il}^m(t, x, \cdot) : [0, T] \times \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ depending explicitly on $\sigma, \nabla_x \sigma, \nabla_\mu \sigma, \psi^m$ and $\mu_{N,t}^{X, \psi^m}$. Based on the signs of $b_{N,il}^m$ and $c_{N,il}^m$, we discretize the terms $b_{N,il}^m(\partial_{x_l} u_i)$ and $c_{N,il}^m(t, x, \cdot)(\partial_{x_l} u_i)(t, \cdot)$ by the upwind finite difference schemes:

$$\begin{aligned} \mathfrak{D}_{N,h,i,l}^{m,j,b}[u(t_j, \cdot)](x) &= b_{N,il}^{m,+}(t_j, x) \frac{u_i(t_j, x + e_l h) - u_i(t_j, x)}{h} + b_{N,il}^{m,-}(t_j, x) \frac{u_i(t_j, x - e_l h) - u_i(t_j, x)}{h}, \\ \mathfrak{D}_{N,h,i,l}^{m,j,c}[u(t_j, \cdot)](x, y) &= c_{N,il}^{m,+}(t_j, x, y) \frac{u_i(t_j, y + e_l h) - u_i(t_j, y)}{h} + c_{N,il}^{m,-}(t_j, x, y) \frac{u_i(t_j, y - e_l h) - u_i(t_j, y)}{h}, \end{aligned}$$

where $\{e_l\}_{l=1}^d \subset \mathbb{R}^d$ is the standard basis of \mathbb{R}^d , and $b^\pm = \max(\pm b, 0)$ for any $b \in \mathbb{R}$. In practice, to evaluate $\mathfrak{D}_{N,h,i,l}^{m,j,c}[u(t_j, \cdot)](x, \cdot)$ on the particles $(X_t^{l,N})_{l=1, \dots, N}$, we shall replace the grid function $(U_k^j)_{k \in \mathbb{Z}^d}$ by its monotone interpolant, which leads to the following approximation:

$$f_{N,h}^{m, \text{ex}}(t_j, x_k, U^j) := \sum_{i,j=1}^d \left(\mathfrak{D}_{N,h,i,l}^{m,j,b}[U^j](x_k) + \mathbb{E}^{\mu_{N,t}^X} [\mathfrak{D}_{N,h,i,l}^{m,j,c}[\mathbf{i}_h[U^j]](x_k, \cdot)] \right). \quad (3.5)$$

Therefore, the fully discrete scheme of (3.1) reads as: $U_k^M = h^m(x_k)$ for all $k \in \mathbb{Z}^d$, and for all $j = 1, \dots, M$, $k \in \mathbb{Z}^d$,

$$U_k^{j-1} - \Delta t L_{N,h}^{m,j-1}[U^{j-1}]_k = U_k^j + \Delta t (f_{N,h}^{m, \text{re}}(t_j, x_k, U^j) + f_{N,h}^{m, \text{ex}}(t_j, x_k, U^j)), \quad (3.6)$$

with $L_{N,h}^{m,j-1}$ defined as in (3.2), $f_{N,h}^{m, \text{re}}$ defined as in (3.3) and $f_{N,h}^{m, \text{ex}}$ defined as in (3.5). As the scheme adopts explicit timestepping for the gradient of u but is implicit in the second order terms, we can set $\Delta t = \mathcal{O}(h)$ for numerical stability.

After obtaining the discrete solution $(U_k^j)_{j,k}$, we follow (2.23) to update the feedback controls ϕ^{m+1} and ψ^{m+1} , which requires us to evaluate $(\nabla F)(\psi^m)(t_j, x_k)$ for all $j = 0, \dots, M$ and $k \in \mathbb{Z}^d$. Observe from (2.18) and (2.14) that f^m and $(\nabla F)(\psi^m)$ have similar structures, except from the fact that f^m depends on $(\partial_x H, \partial_\mu H)$, and $(\nabla F)(\psi^m)$ depends on $(\partial_a H, \partial_\nu H)$. Hence, one can construct an analogue approximation of $(\nabla F)(\psi^m)(t_j, x_k)$ by replacing $(\partial_x b, \partial_\mu b, \partial_x \sigma, \partial_\mu \sigma)$ in $f_N^{m, \text{re}} + f_N^{m, \text{ex}}$ with $(\partial_a b, \partial_\nu b, \partial_a \sigma, \partial_\nu \sigma)$. This enables us to evaluate ϕ^{m+1} and ψ^{m+1} at all grid points (t_j, x_k) , and subsequently to obtain a particle approximation of X^{m+1} for the next NAG iteration based on (2.24) and the monotone interpolant of ψ^{m+1} over the grids.

⁴Here δ_{kj} is the Kronecker delta.

4 Implementation of the FIPDE method via residual approximation

Despite the finite difference approximation in Section 3 being very effective in solving low-dimensional MFC problems, it cannot be applied to MFC problems with high-dimensional state processes, due to unaffordable computational costs. In this section, we shall propose a neural network-based implementation of the FIPDE method, where at each NAG iteration, we compute an approximate solution to (2.15) by minimizing a proper residual over a family of neural networks.

As in Section 3, we focus on the m -th NAG iteration and seek a vector-valued function $u : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ satisfying the following nonlocal parabolic PDE system:

$$(\partial_t u)(t, x) + (\mathcal{L}_N^m u)(t, x) = -f_N^m(t, x, u, v), \quad (4.1)$$

with $u(T, x) = h^m(x)$ and $v(t, x) = (\nabla_x u)(t, x)\sigma^m(t, x)$ for all $(t, x) \in [0, T] \times \mathbb{R}^d$. The operator \mathcal{L}_N^m and the function f_N^m (which is affine and nonlocal in u, v) are defined as in (3.1), which depend on a given feedback control ψ^m (represented by a multilayer neural network) and the empirical measures of a given particle approximation $(X^{l,N})_{l=1,\dots,N}$ of the state process X^m for the present NAG iteration.

In the following, we shall reformulate (4.1) into an empirical risk minimization problem over multilayer neural networks, which is then solved by using stochastic gradient descent (SGD) algorithms; see e.g., the Deep Galerkin Method (DGM) in [52]. However, instead of directly applying DGM to (4.1), we shall consider an acceleration method by first decomposing the solution u into:

$$u(t, x) = \bar{u}(t, x) + \tilde{u}(t, x), \quad (t, x) \in [0, T] \times \mathbb{R}^d, \quad (4.2)$$

where \bar{u} is an approximate solution to (4.1) computed by some efficient numerical methods, and \tilde{u} is a residual correction of \bar{u} based on neural networks. By computing the dominant part of the solution u efficiently and merely applying SGD algorithms for the small residual term, we can obtain an accurate and stable approximation of u with a small number of SGD iterations, and subsequently reduce the total computation time for solving PDE systems at all NAG iterations.

In general, the numerical solver of \bar{u} should be designed in a problem dependent way. For many practical MFC problems (see, e.g., [31, 5]), we can first linearize the dynamics (1.1) around the target states and approximate the cost functions (1.2) by suitable quadratic costs. This leads to a linear-quadratic (LQ) approximation of the MFC problem (1.1)-(1.2), and the approximate solution \bar{u} of (4.1) can be chosen as the solution of the (matrix-valued) differential Riccati equations for the resulting LQ MFC problem. Consequently, the decomposition (4.2) can be viewed as a neural network-based nonlinear correction to the (suboptimal) linear feedback control. We refer the reader to Section 5.2.2 for more details on the LQ approximation of MFC problems with nonsmooth costs.

Given the approximate solution $\bar{u} : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$, we see from (4.1) that the residual term $\tilde{u} : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ satisfies for all $(t, x) \in [0, T] \times \mathbb{R}^d$,

$$(\partial_t \tilde{u})(t, x) + (\mathcal{L}_N^m \tilde{u})(t, x) = -\tilde{f}_N^m(t, x, \tilde{u}, \tilde{v}), \quad \tilde{u}(T, x) = h^m(x) - \bar{u}(T, x), \quad (4.3)$$

where $\tilde{v}(t, x) = (\nabla_x \tilde{u})(t, x)\sigma^m(t, x)$, and

$$\tilde{f}_N^m(t, x, \tilde{u}, \tilde{v}) = f_N^m(t, x, \tilde{u} + \bar{u}, \tilde{u} + \tilde{v}) + (\partial_t \bar{u})(t, x) + (\mathcal{L}_N^m \bar{u})(t, x).$$

We then extend the residual based method for scalar PDEs in [52, 34] to the coupled PDE system (4.3). In particular, let $\mathcal{D} \subset \mathbb{R}^d$ be the chosen computational domain and $\mathcal{N}_u = \{u^\theta : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d \mid \theta \in \mathbb{R}^p\}$ be a family of multilayer neural networks with some prescribed architectures and sufficiently smooth activation functions. Then we seek the optimal neural network in \mathcal{N}_u to approximate \tilde{u} by minimising the following loss function over the weights θ :

$$\begin{aligned} \mathcal{E}(\theta) = & \|(\partial_t \tilde{u}^\theta)(\cdot, \cdot) + (\mathcal{L}_N^m \tilde{u}^\theta)(\cdot, \cdot) + \tilde{f}_N^m(\cdot, \cdot, \tilde{u}^\theta, (\nabla_x \tilde{u}^\theta)\sigma^m)\|_{[0, T] \times \mathcal{D}, \nu_1}^2 \\ & + \eta_1 \|\tilde{u}^\theta(T, \cdot) - (h^m(\cdot) - \bar{u}(T, \cdot))\|_{\mathcal{D}, \nu_2}^2 + \eta_2 \|\tilde{u}^\theta(\cdot, \cdot)\|_{[0, T] \times \partial \mathcal{D}, \nu_3}^2. \end{aligned} \quad (4.4)$$

Here, $\nu_i, i = 1, 2, 3$, are some given probability measures on $[0, T] \times \mathcal{D}$, \mathcal{D} and $[0, T] \times \partial \mathcal{D}$ respectively, and $\eta_1, \eta_2 > 0$ are some given weighting parameters (possibly different among all NAG iterations) introduced

to balance the interior residual and the residuals of the boundary data. Note that \tilde{u}^θ takes values in \mathbb{R}^d , and (4.4) contains residuals of all components.

In practice, the loss function (4.4) can be minimized by using SGD based on a sequence of mini-batches of pseudorandom points or quasi-Monte Carlo points. More precisely, for the j -th SGD iteration with $j \in \mathbb{N}$, we first generate N_{in} points from $[0, T] \times \mathcal{D}$, N_{ter} points from \mathcal{D} and N_{bdy} points from $[0, T] \times \partial\mathcal{D}$ according to the measures ν_1 , ν_2 and ν_3 , respectively, then evaluate the following empirical loss with the current weights $\theta_j \in \mathbb{R}^k$:

$$\begin{aligned} \mathcal{E}_{\text{em}}(\theta_j) = & \frac{1}{N_{\text{in}}} \sum_{i=1}^{N_{\text{in}}} |(\partial_t \tilde{u}^{\theta_j})(t_i, x_i) + (\mathcal{L}_N^m \tilde{u}^{\theta_j})(t_i, x_i) + \tilde{f}_N^m(t_i, x_i, \tilde{u}^{\theta_j}, (\nabla_x \tilde{u}^{\theta_j}) \sigma^m)|^2 \\ & + \frac{\eta_1}{N_{\text{ter}}} \sum_{i=1}^{N_{\text{ter}}} |\tilde{u}^{\theta_j}(T, x_i) - (h^m(x_i) - \bar{u}(T, x_i))|^2 + \frac{\eta_2}{N_{\text{bdy}}} \sum_{i=1}^{N_{\text{bdy}}} |\tilde{u}^{\theta_j}(t_i, x_i)|^2, \end{aligned} \quad (4.5)$$

and finally obtain the updated weights $\theta_{j+1} = \theta_j - \tau_j (\nabla_{\theta} \mathcal{E}_{\text{em}})(\theta_j)$ with a stepsize $\tau_j > 0$. The above SGD iterations are performed until the index j meets the maximum iteration number or an accuracy tolerance is satisfied.

Finally, we discuss the implementation of Step 5 in Algorithm 2 based on an approximate solution \tilde{u}^θ to (4.3) and the current feedback control ψ^m , where both are represented by a multilayer neural network. By (4.2), the solution u^m to (4.1) is now approximated by $u^\theta := \bar{u} + \tilde{u}^\theta$, which leads to a pointwise approximation $(\nabla F)(\psi^m, \theta)$ of the function $(\nabla F)(\psi^m)$ by replacing u^m in (2.14) with u^θ . Note that the required derivative $\nabla_x u^\theta$ for approximating the function v^m can be computed analytically if \bar{u} and \tilde{u}^θ are differentiable. However, in contrast to the grid-based representation of feedback controls in Section 3, the neural network representation of the feedback control prevents us from obtaining the updated control ψ^{m+1} in (2.23) via simple operations applied to the parameters. This is due to the fact that a linear combination of multilayer neural networks with nonlinear activation functions in general can only be expressed as a neural network with more complicated architectures (see e.g., [50, Lemma A.1]). Hence, exactly following the update rules in (2.23) would require us to save all $(\psi^m)_{m \in \mathbb{N} \cup \{0\}}$ and $(u^m)_{m \in \mathbb{N} \cup \{0\}}$, which increases both the memory requirements and the computational cost for Steps 3 and 4 of Algorithm 2 if the NAG iteration m is large. To overcome this difficulty, we shall follow (2.23) approximately and obtain the updated controls as follows:

$$\begin{aligned} \phi^{m+1} & \in \arg \min_{\phi \in \mathcal{N}_\phi} \|\phi - (\text{prox}_{\tau\ell}(\psi^m - \tau(\nabla F)(\psi^m, \theta)))\|_{[0, T] \times \mathcal{D}, \nu_1}^2, \\ \psi^{m+1} & \in \arg \min_{\psi \in \mathcal{N}_\psi} \|\psi - (\phi^{m+1} + \frac{m}{m+3}(\phi^{m+1} - \phi^m))\|_{[0, T] \times \mathcal{D}, \nu_1}^2, \end{aligned} \quad (4.6)$$

where \mathcal{N}_ϕ and \mathcal{N}_ψ are families of multilayer neural networks with prescribed architectures and unknown parameters. These supervised learning problems can be easily solved by performing gradient descent based on sample points from $[0, T] \times \mathcal{D}$.

5 Numerical experiments

In this section, we demonstrate the effectiveness of the FIPDE scheme through numerical experiments. We present two MFC problems with nonsmooth optimization objectives: a portfolio liquidation problem with trade crowding and transaction costs in Section 5.1, and sparse consensus control of two-dimensional and six-dimensional stochastic Cucker–Smale models in Section 5.2. For both examples, we benchmark the FIPDE method with existing pure data-driven algorithms, including the empirical regression method in [38] and the neural network-based policy gradient method in [21], which lead to approximate feedback control merely along trajectories of the optimal state process. Our experiments show that compared with pure data-driven approaches, the FIPDE method leads to a global approximation of the optimal feedback control, and achieves a robust performance in terms of model perturbations.

5.1 Portfolio liquidation with trade crowding and transaction costs

In this section, we consider a portfolio liquidation problem where a large number of market participants try to liquidate their positions on the same asset by a given terminal time $T > 0$ (see [12, 1]), while taking

into account (possibly nonsmooth) execution costs and the permanent price impact caused by their trading actions. The cooperative equilibrium leads to a MFC problem for a representative agent.

Let $(\alpha_t)_{t \in [0, T]}$ be the trading speed chosen by the representative agent, the state dynamics of the MFC problem is given by: for all $t \in [0, T]$,

$$dQ_t = \alpha_t dt, \quad dS_t = \lambda \mathbb{E}[\alpha_t] dt + \sigma dW_t,$$

where $Q = (Q_t)_{t \in [0, T]}$ is the inventory process with a *random* initial state Q_0 representing the initial inventories for all participants, $S = (S_t)_{t \in [0, T]}$ is the price process with a deterministic initial state $S_0 = s_0 \geq 0$, and $\lambda \mathbb{E}[\alpha_t]$ with $\lambda \geq 0$ represents the permanent market impact on the asset price due to the trading of all participants. Here, the processes Q , S and W are all one-dimensional.

The objective of the agent is then to minimize the following cost functional

$$J(\alpha; (S_0, Q_0)) = \mathbb{E} \left[\int_0^T \left(\alpha_t S_t^\alpha + (Q_t^\alpha)^2 + k_1 |\alpha_t|^2 + k_2 |\alpha_t| \right) dt - Q_T^\alpha (S_T^\alpha - \gamma Q_T^\alpha) \right], \quad (5.1)$$

over all possible trading speeds α , where $(Q_t^\alpha)^2$ penalizes the current inventory, $k_1, k_2 \geq 0$, and $Q_T^\alpha (S_T^\alpha - \gamma Q_T^\alpha)$ with $\gamma \geq 0$ is the liquidation value of the remaining inventory at terminal time. Note that the nonsmooth term $k_2 |\alpha|$ models proportional execution costs such as the bid-ask spread, the fees paid to the venue, and/or a stamp duty (see, e.g., [32]). In the following, we shall perform experiments with different choices of Q_0 and k_2 (whose values will be specified later) while fixing the other model parameters as: $T = 1$, $s_0 = 2$, $\lambda = 0.5$, $\sigma = 0.7$, $\gamma = 0.5$ and $k_1 = 1$.

We initialize Algorithm 2 with stepsize $\tau = 1/6$ and initial guess $\phi^0 = 0$. At the m -th NAG iteration, given an approximate control strategy $\psi^m : [0, T] \times \mathbb{R}^2 \rightarrow \mathbb{R}$ with associated state processes (S^m, Q^m) , we consider the following decoupled system of PDEs (cf., (2.15)): for all $(t, s, q) \in [0, T] \times \mathbb{R} \times \mathbb{R}$,

$$(\partial_t u_1)(t, s, q) + (\mathcal{L}^m u_1)(t, s, q) = -\psi^m(t, s, q), \quad u_1(T, s, q) = -q, \quad (5.2a)$$

$$(\partial_t u_2)(t, s, q) + (\mathcal{L}^m u_2)(t, s, q) = -2q, \quad u_2(T, s, q) = -s + 2\gamma q, \quad (5.2b)$$

with \mathcal{L}^m such that for each $\varphi \in C^{1,2}([0, T] \times \mathbb{R}^2; \mathbb{R})$ and $(t, s, q) \in [0, T] \times \mathbb{R} \times \mathbb{R}$,

$$(\mathcal{L}^m \varphi)(t, s, q) = \frac{1}{2} \sigma^2 (\partial_{ss} \varphi)(t, s, q) + \lambda \mathbb{E}[\psi^m(t, S_t^m, Q_t^m)] (\partial_s \varphi)(t, s, q) + \psi^m(t, s, q) (\partial_q \varphi)(t, s, q). \quad (5.3)$$

We implement the FIPDE algorithm with the monotone scheme (3.6) to solve the system (5.2). We first approximate the expectations in (5.3) by empirical averages over $N = 10^4$ particle approximations of (S^m, Q^m) generated by (2.24) with $\Delta t = 1/50$, localize the equation on the domain $\mathcal{D} = [-2, 6] \times [0, 4]$ and impose boundary conditions as the terminal condition, i.e., $(u_1(t, s, q), u_2(t, s, q)) = (-q, -s + 2\gamma q)$ for all $(t, s, q) \in [0, T] \times \partial \mathcal{D}$. Our experiments with larger computational domains indicate that this domain truncation and boundary condition leads to a negligible domain truncation error. Then we construct an implicit first-order monotone scheme (3.6) for the localized system (5.2) by discretizing the first-order and second-order derivatives in (5.3) via the upwind finite difference and the central difference, respectively. The chosen time stepsize is $\Delta t = 1/50$, and the spatial mesh sizes $h_s = 8/50$, $h_q = 4/50$.

As a benchmark for the FIPDE scheme, we also implement an empirical regression (EMReg) method to solve (5.1), where for each NAG iteration, we solve the adjoint BSDE (2.4) by projecting the decoupling fields (u^m, v^m) of (Y^m, Z^m) on prescribed vector spaces of basis functions, and evaluating the coefficients by performing regressions based on simulated trajectories of (S^m, Q^m) (see, e.g., [38]). In particular, we partition the computation domain $\mathcal{D} = [-2, 6] \times [0, 4]$ with the same meshsize $h_s = 8/50$ and $h_q = 4/50$ as that of the FIPDE method, and choose the indicator functions of all subcells as the basis functions in the regression. We also employ (2.24) with the same parameters $N = 10^4$ and $M = 50$ as those for the FIPDE method to generate state trajectories, in order to ensure a fair comparison. Other implementation details are given in Appendix A.1.

We first examine the performance of the FIPDE and EMReg schemes for solving the linear-quadratic (LQ) MFC problem (5.1) with $Q_0 \sim \mathcal{U}(1, 2)$ (the uniform distribution on $(1, 2)$) and $k_2 = 0$. Extending the arguments in [54] to (5.1) yields that for any square-integrable initial condition Q_0 , the optimal feedback control of (5.1) is of the form:

$$\phi^*(t, s, q) = a_t q + b_t \mathbb{E}[Q_t^*], \quad \forall (t, s, q) \in [0, T] \times \mathbb{R} \times \mathbb{R}, \quad (5.4)$$

where Q^* is the optimal inventory process, and $a_t, b_t : [0, T] \rightarrow \mathbb{R}$ are solutions to some Riccati equations depending explicitly on T, λ, γ and k_1 , but independent of s_0 and Q_0 . Despite the fact that both the FIPDE and EMReg schemes achieve less than 1% absolute error for value function approximations within 5 NAG iterations (see Figure 11 in Appendix A.1 for more details), these two methods generate qualitatively different feedback controls. Figure 1 compares the exact feedback control (5.4) and the approximate feedback controls obtained by the FIPDE and EMReg schemes after 20 NAG iterations, for which we evaluate the feedback strategies at $s = 2$ and $(t, q) \in [0, 1] \times [0.5, 2.5]$. One can clearly observe that the approximate feedback control from the FIPDE scheme is in almost exact agreement with the analytic solution on the entire computational domain. In contrast, the EMReg method produces a much more irregular control strategy, which is only accurate on a certain part of the domain. Recall that we initialize the NAG iteration in Algorithm 2 with $\phi^0 = 0$, and Figure 1 (right) indicates that the EMReg method does not update the initial guess for $q > 2$. In fact, as the EMReg method approximates the adjoint processes by performing regression based on the simulated trajectories of the state process, the EMReg method not only suffers more from statistical errors, but also cannot recover the exact feedback control beyond the support of the simulated trajectories. As we shall see soon, such a local approximation property makes these pure data-driven approaches unstable with respect to model perturbation.




Figure 1: Feedback controls of the LQ optimal liquidation problem with $Q_0 \sim \mathcal{U}(1, 2)$ obtained by using different methods.

We proceed to investigate the robustness of feedback controls obtained by the FIPDE scheme and the EMReg scheme with respect to perturbation of the random initial condition Q_0 , which models the initial inventories of all market participants. As it is practically difficult to obtain the precise distribution of Q_0 , it is desirable for the feedback control to be stable in terms of the uncertainty in the initial law (see, e.g., [27]). In particular, we first obtain approximate feedback controls by applying the FIPDE and EMReg schemes to solve (5.1) with $Q_0^{\text{pre}} \sim \mathcal{U}(1, 2)$, and then examine the performance of these feedback controls on perturbed models with $Q_0 \sim \mathcal{U}(Q_{0,\min}, Q_{0,\max})$ for $Q_{0,\min} \in [0.5, 1.5]$ and $Q_{0,\max} \in [1.5, 2.5]$.

Figure 2 illustrates the performance of these precomputed controls on models with different Q_0 in terms of the absolute performance gap $|J(\phi^{\text{pre}}) - J^*|$, where $J(\phi^{\text{pre}})$ is the expected cost of a precomputed feedback control ϕ^{pre} on the perturbed model, and J^* is the optimal cost of the perturbed model. One can clearly observe from Figure 2 (left) that the precomputed feedback control from the FIPDE scheme is very robust with respect to perturbation of Q_0 , as it yields extremely small performance gaps for all perturbed models. Moreover, the absolute performance gaps remain almost constants for different Q_0 with the same expectation. This is because the feedback control from the FIPDE scheme captures the precise spatial dependence of the optimal feedback control of a perturbed problem (i.e., the function ϕ^* in (5.4)), while it keeps the measure dependence the same as that for the unperturbed model. As the optimal feedback control of the LQ MFC problem (5.1) depends on the law of Q_0 only through its first moment, the absolute performance gap $|J(\phi^{\text{pre}}) - J^*|$ for the FIPDE scheme is purely determined by the perturbation of $\mathbb{E}[Q_0]$.

In contrast, as shown in Figure 2 (right), the approximate feedback control from the EMReg scheme is very sensitive to perturbations of Q_0 , where absolute performance gaps are typically a few magnitudes larger than those of the FIPDE scheme. This phenomenon is more pronounced if the support of the perturbed Q_0 is not contained by the interval $[1, 2]$, i.e., the support of the original Q_0 . This is because the EMReg scheme (and other pure data-driven algorithms) computes approximate feedback controls

merely based on trajectories of the original state process, and consequently the resulting feedback control performs poorly along the trajectories outside the support of the original system. In particular, since it is critical to execute a proper strategy for large values of Q_t in the present liquidation problem, a slight perturbation of $Q_{0,\max}$ will significantly worsen the performance of the precomputed control.

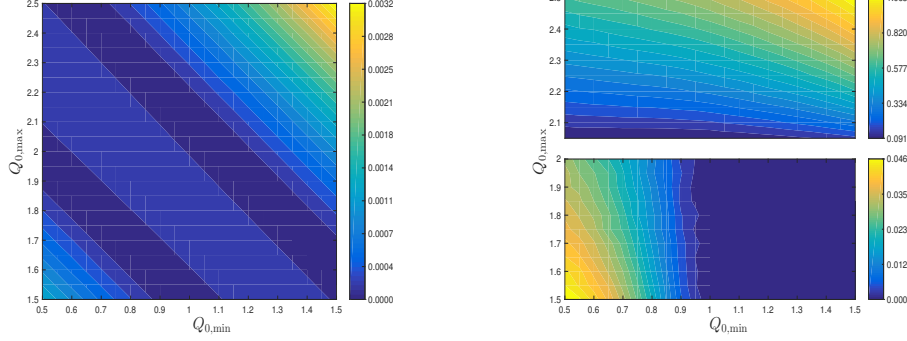


Figure 2: Absolute performance gaps of precomputed feedback controls from the FIPDE scheme (left) and the EMReg scheme (right) on perturbed models with $Q_0 \sim \mathcal{U}(Q_{0,\min}, Q_{0,\max})$.

The improved robustness of the FIPDE scheme over the EMReg scheme can be better analyzed by the relative performance gaps $|J(\phi^{\text{pre}}) - J^*|/|J^*|$, whose distributions (with fixed $Q_{0,\min}$ and varying $Q_{0,\max}$) are summarized by the box plots in Figure 3. We can clearly see that the precomputed control of the FIPDE scheme achieves a relative error of less than 2% on most perturbed models (Figure 2 (left)), while the precomputed control of the EMReg scheme will typically lead to a relative error ranging from 20% to $10^3\%$ (Figure 3 (top-right)).⁵

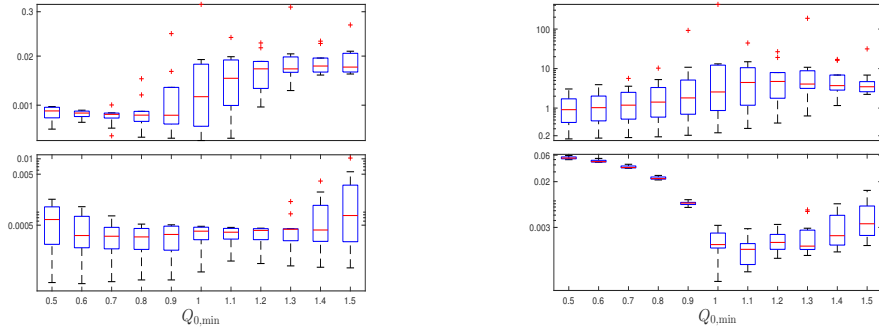


Figure 3: Relative performance gaps of precomputed feedback controls on perturbed models with $Q_0 \sim \mathcal{U}(Q_{0,\min}, Q_{0,\max})$; from left to right: results for the FIPDE and EMReg scheme; from top to bottom: distributions of relative errors for fixed $Q_{0,\min}$ and varying $Q_{0,\max} \in [2, 2.5]$ and $Q_{0,\max} \in [1.5, 2]$, where on each box, the central line is the median, the edges of the box are the 25th and 75th percentiles, the whiskers extend to non-outliers extreme data points, and outliers are plotted individually.

Finally, we examine the performance of the FIPDE scheme for solving the nonsmooth MFC problem (5.1) with $k_1 = k_2 = 1$ in (5.1). We initialize Algorithm 2 with $\tau = 1/6$ and $\phi^0 = 0$, and for each NAG iteration, discretize the PDE system (5.2) by using the same monotone scheme as that for the LQ case. Note that for all $\tau > 0$, the proximal operator of $\tau|\cdot|$ is given by $\text{prox}_{\tau|\cdot|}(x) = \text{sgn}(x)(|x| - k_2\tau)$ for all $|x| \geq k_2\tau$ and $\text{prox}_{\tau\ell}(x) = 0$ for all $|x| \leq k_2\tau$.

We carry out the FIPDE scheme with 20 NAG iterations, whose convergence is shown in Figure 11 in Appendix A.1; note that in fact 4 NAG iterations are sufficient to approximate the value function accurately. The resulting approximate feedback control is independent of the variable s but nonlinear in the variable q . Figure 4 (left) presents the approximate feedback controls for $s = 2$ and $(t, q) \in$

⁵ We have ignored the outliers (marked as plus signs) in Figure 3 (top), which resulted from evaluating relative errors with optimal costs J^* very close to zero.

$[0, 1] \times [0.5, 2.5]$, which clearly shows that the nonsmooth cost enhances the sparsity of the optimal strategy, especially near the terminal time. We then analyze the robustness of the FIPDE scheme by exercising the feedback control ϕ^{pre} computed with $Q_0 \sim \mathcal{U}(1, 2)$ on perturbed models with $Q_0 \sim \mathcal{U}(Q_{0,\min}, Q_{0,\max})$. Due to the absence of analytic solution for the nonsmooth MFC problem, we compare the expected cost $J(\phi^{\text{pre}})$ of ϕ^{pre} on the perturbed model against the numerical approximation $J(\phi^{\text{pert}})$ of the optimal cost of the perturbed model, where ϕ^{pert} is the feedback control obtained by the FIPDE scheme with the perturbed Q_0 . Figure 4 (right) depicts the absolute performance gaps $|J(\phi^{\text{pre}}) - J(\phi^{\text{pert}})|$ for different Q_0 . Due to the nonsmooth cost function, the optimal feedback controls of (5.1) depend nonlinearly on the law of Q_0 , and hence the performance gaps are no longer constant along the diagonals (cf., Figure 2 (left)). However, the absolute performance gaps remain small for all perturbations, which demonstrates the robustness of the FIPDE scheme in the present nonsmooth setting.

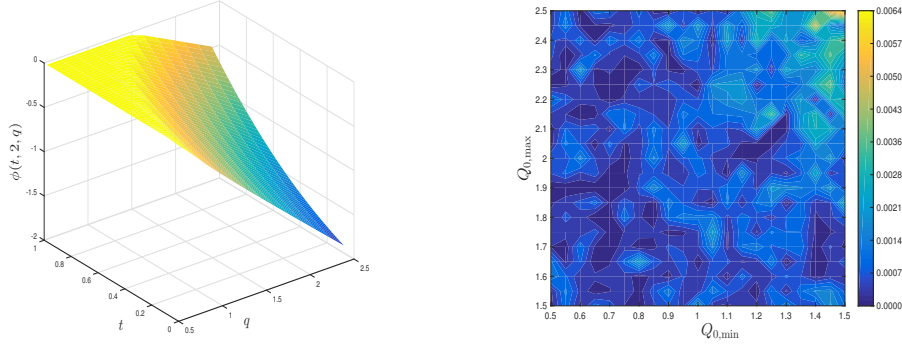


Figure 4: Numerical results of the FIPDE scheme for the nonsmooth optimal liquidation problem with $k_2 = 1$; from left to right: approximate feedback control for $Q_0 \sim \mathcal{U}(1, 2)$, and absolute performance gaps of the precomputed feedback control on perturbed models with $Q_0 \sim \mathcal{U}(Q_{0,\min}, Q_{0,\max})$.

5.2 Sparse consensus control of stochastic Cucker–Smale models

In this section, we study optimal control of the multidimensional stochastic mean-field Cucker–Smale (C-S) dynamics (see, e.g., [43] and [20, Chapter 4]), where the controller aims to enforce consensus emergence of an interactive particle system via (possibly sparse) external intervention. In its general form this leads to a nonsmooth nonconvex MFC problem.

Given a terminal time $T > 0$ and a d -dimensional adapted control strategy $(\alpha_t)_{t \in [0, T]}$, we consider the following $2d$ -dimensional controlled dynamics with additive noise, which can be viewed as the large population limit of the finite-particle model studied in [24]: for all $t \in [0, T]$,

$$dx_t = v_t dt, \quad dv_t = \left(\int_{\mathbb{R}^d \times \mathbb{R}^d} \kappa(x_t, v_t, x', v') \mathcal{L}_{(x_t, v_t)}(dx', dv') + \alpha_t \right) dt + \sigma dW_t, \quad (5.5)$$

with initial state $(x_0, v_0) \in L^2(\Omega; \mathbb{R}^d \times \mathbb{R}^d)$, where W is a d -dimensional Brownian motion defined on a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \in [0, T]}, \mathbb{P})$, $\sigma \in \mathbb{R}^{d \times d}$ and $\kappa : \mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ is given by

$$\kappa(x, v, x', v') = \frac{K(v' - v)}{(1 + |x - x'|^2)^\beta}, \quad \text{with some } \beta, K \geq 0.$$

It is well-known that for uncontrolled deterministic models (with $\sigma = 0$ and $\alpha \equiv 0$), a time-asymptotic flocking behaviour (i.e., all trajectories of the velocity process tend to the same value as $t \rightarrow \infty$) only appears for $\beta \leq 1/2$ or for specific initial conditions if $\beta > 1/2$; see [24, 23, 9] and references therein. Moreover, as shown in [24, 43], even when the deterministic counterpart exhibits a time-asymptotic flocking behaviour, the additive noise may prevent the emergence of flocking in the stochastic model.

The aim of the controller is to either induce consensus on models that would otherwise diverge, or to accelerate the flocking for an initial configuration that would naturally self-organise. More precisely, for given constants $\gamma_1, \gamma_2 \geq 0$, we consider minimizing the following cost functional

$$J(\alpha; (x_0, v_0)) = \mathbb{E} \left[\int_0^T \left(|v_t - \mathbb{E}[v_t]|^2 + \gamma_1 |\alpha_t|^2 + \gamma_2 |\alpha_{t,1}| \right) dt + |v_T - \mathbb{E}[v_T]|^2 \right] \quad (5.6)$$

over all adapted control processes α taking values in \mathbb{R}^d , where $\|\cdot\|_1$ is the ℓ_1 -norm of a given vector. Note that in general, the cost functional J is neither convex nor smooth in the control process α , due to the nonlinear interaction kernel κ in (5.5) and the ℓ_1 -norm in (5.6). However, in the special case with $\beta = \gamma_2 = 0$ and $\gamma_1 > 0$, (5.5)-(5.6) is a LQ MFC problem whose optimal feedback control can be found via Riccati equations (see, e.g., [54, 12]).

In the subsequent two sections, we demonstrate the effectiveness of the FIPDE method (i.e., Algorithm 2) to solve (5.5)-(5.6) with different choices of β, d and γ_2 . For the m -th NAG iteration, given the approximate feedback control $\psi^m : [0, T] \times \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ with associated state processes $(x_t^m, v_t^m)_{t \in [0, T]}$, the FIPDE method seeks functions $u_1, u_2 : [0, T] \times \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ satisfying the following coupled system of $2d$ -dimensional parabolic PDEs: for all $(t, x, v) \in [0, T] \times \mathbb{R}^d \times \mathbb{R}^d$,

$$(\partial_t u_1 + \mathcal{L}^m u_1)(t, x, v) = -f_1^m(t, x, v, u_2), \quad u_1(T, x, v) = 0, \quad (5.7a)$$

$$(\partial_t u_2 + \mathcal{L}^m u_2)(t, x, v) = -f_2^m(t, x, v, u_1, u_2), \quad u_2(T, x, v) = 2(v - \mathbb{E}[v_T^m]), \quad (5.7b)$$

with the operator \mathcal{L}^m and the source terms f_1^m, f_2^m satisfying for all $\varphi \in C^{1,2}([0, T] \times \mathbb{R}^{2d}; \mathbb{R})$,

$$\begin{aligned} (\mathcal{L}^m \varphi)(t, x, v) &= \frac{1}{2} \text{tr}(\sigma \sigma^\top (\partial_{vv} \varphi)(t, x, v)) + v^\top (\partial_x \varphi)(t, x, v) \\ &\quad + (\mathbb{E}[\kappa(x, v, x_t^m, v_t^m)] + \psi^m(t, x, v))^\top (\partial_v \varphi)(t, x, v), \\ f_1^m(t, x, v, u_2) &= \mathbb{E}[(\partial_x \kappa)(x, v, x_t^m, v_t^m)] u_2(t, x, v) + \mathbb{E}[(\partial_{x'} \kappa)(x_t^m, v_t^m, x, v) u_2(t, x_t^m, v_t^m)], \\ f_2^m(t, x, v, u_1, u_2) &= u_1(t, x, v) + 2(v - \mathbb{E}[v_t^m]) + \mathbb{E}[(\partial_v \kappa)(x, v, x_t^m, v_t^m)] u_2(t, x, v) \\ &\quad + \mathbb{E}[(\partial_{v'} \kappa)(x_t^m, v_t^m, x, v) u_2(t, x_t^m, v_t^m)], \end{aligned} \quad (5.8)$$

where for each $w \in \{x, x', v, v'\}$, $\partial_w \kappa$ denotes the Jacobian matrix such that $(\partial_w \kappa)_{ij} = \partial_{w_i} \kappa_j$. In practice, we approximate the expectations in the coefficients of (5.8) by the corresponding empirical averages over particle approximations of (x^m, v^m) , which are generated by (2.24) with sufficiently large $N, M \in \mathbb{N}$. The PDE systems with approximated coefficients will be solved by using the finite difference approximation in Section 3 or the residual approximation method in Section 4, depending on the problem dimension $2d$; see Sections 5.2.1 and 5.2.2, respectively, for more details.

We shall also compare the FIPDE method with the neural network-based policy gradient (NNPG) method, which is a pure data-driven algorithm proposed in [21] to solve MFC problems. The NNPG method considers minimizing (5.6) over a class of feedback controls represented by multilayer neural networks with weights θ , and obtains the optimal weights by applying gradient descent algorithms based on simulated trajectories of the state process (5.5) (see Sections 5.2.1 and 5.2.2 for more details). As we shall see soon, the proposed FIPDE method achieves more accurate and interpretable feedback controls than the NNPG method, in both the low-dimensional and high-dimensional settings.

5.2.1 Consensus control of two-dimensional C-S models

This section studies (5.5)-(5.6) for two-dimensional C-S models with different communication rates β . In particular, we shall perform experiments with the following model parameters: $T = d = K = 1$, $\sigma = \gamma_1 = 0.1$, $\gamma_2 = 0$, $\beta \in \{0, 10\}$, and (x_0, v_0) follows the two-component Gaussian mixture distribution with mixture weights $(0.5, 0.5)$, mean $(1.2, 1.8)$ and covariance $0.01\mathbb{I}_2$ for component 1, and mean $(1.8, 1.2)$ and covariance $0.01\mathbb{I}_2$ for component 2, where \mathbb{I}_2 is the 2×2 identity matrix.

We now discuss the implementation details of the FIPDE algorithm with finite difference approximation for the present two-dimensional setting. The FIPDE algorithm is initialized with $\tau = 1/6$ and $\phi^0 \equiv 0$. At the m -th NAG iteration, given an approximate feedback control ψ^m , we first generate particle approximations $(x^{m,l}, v^{m,l})_{l=1, \dots, 10^4}$ of the state process (x^m, v^m) with time stepsize $\Delta t = 1/50$ (i.e., (2.24) with $N = 10^4$ and $M = 50$), and replace the expectations in (5.8) by the empirical averages over particles. The PDE system with approximate coefficients is then localized on the domain $\mathcal{D} = [0, 5] \times [0, 4]$ with boundary conditions being the terminal values, i.e., $(u_1(t, x, v), u_2(t, x, v)) = (0, 2(v - \frac{1}{N} \sum_{l=1}^N v_T^{m,l}))$ for all $(t, x, v) \in [0, T] \times \partial\mathcal{D}$. We further construct a semi-implicit first-order monotone scheme (3.6) for the localized system (5.7) by adopting the time stepsize $\Delta t = 1/50$ and mesh sizes $h_x = 5/50$, $h_v = 4/50$, and discretizing the first-order derivatives in (5.8) via the upwind finite difference and the second-order derivative in (5.8) via the central difference.

For the LQ MFC problem (5.5)-(5.6) with $\beta = 0$, we compare the approximate feedback control from the FIPDE scheme with the optimal feedback control given by

$$\phi^*(t, x, v) = -\frac{a_t}{2\gamma_1}(v - \mathbb{E}[v_t^*]), \quad (t, x, v) \in [0, T] \times \mathbb{R} \times \mathbb{R}, \quad (5.9)$$

where $a : [0, T] \times \mathbb{R}$ satisfies $a'_t - 2Ka_t - \frac{1}{2\gamma_1}a_t^2 + 2 = 0$ with $a_T = 2$, and v^* is the optimal velocity process (see, e.g., [54, 12]). The blue line in Figure 5 (left) presents the expected costs of the approximate feedback controls obtained by the FIPDE method at all NAG iterations, which converge exponentially to the optimal cost in terms of the number of NAG iterations; a linear regression of the data shows the absolute error for the m -th iteration is of the magnitude $\mathcal{O}(0.55^m)$. The FIPDE method is benchmarked against the iterative PDE-based (IPDE) algorithm Algorithm 1 (i.e., $\psi^m = \phi^m$ for all m), whose convergence is shown by the red line in Figure 5 (left). Note that the FIPDE and IPDE methods take a similar time to perform one gradient descent iteration, but the FIPDE method requires much fewer iterations to achieve high accuracy. This shows that the momentum step indeed accelerates the algorithm convergence, and justifies the terminology “fast” in FIPDE. We further depict the approximate feedback control ϕ (generated by the last NAG iteration) with $t = 0$ and $(x, v) \in [1, 2]^2$ in Figure 5 (right), which approximates the optimal feedback control $\phi^*(0, \cdot, \cdot)$ (see (5.9)) with a relative error of 1.7% in the L^2 -norm. One can clearly observe that the approximate feedback control captures the affine structure of ϕ^* in v , constant in x .

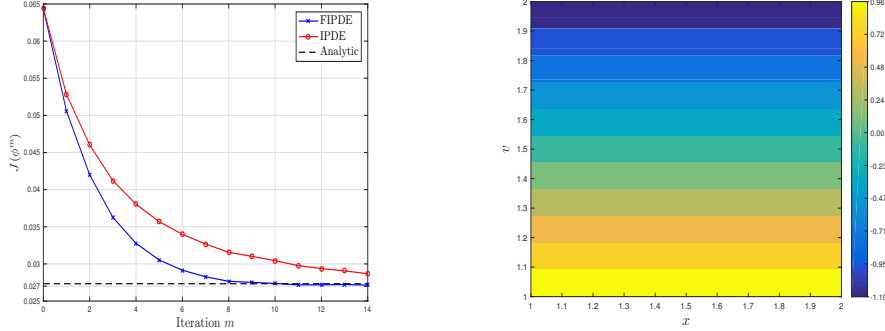


Figure 5: Numerical results of the FIPDE scheme for two-dimensional C-S model with $\beta = 0$; from left to right: convergence of the FIPDE method in terms of NAG iterations, and the approximate feedback control at $t = 0$.

Then we proceed to study the (nonconvex) MFC problem (5.5)-(5.6) with $\beta = 10$, where the uncontrolled velocity process does not exhibit a flocking behaviour; see Figure 12 in Appendix A.2. In the following, we apply the FIPDE method (with the same discretization parameters as above) to design feedback controls and induce consensus for such models. We shall further benchmark the performance of the FIPDE method against a neural network-based policy gradient (NNPG) method (see Algorithm 1 in [21]), whose implementation details are given in Appendix A.2. In particular, we seek an approximate feedback control among a family of neural networks whose flexibility is sufficient to capture the nonlinearity of the optimal feedback control in (t, x, v) , and obtain the optimal neural network approximation by running the Adam algorithm [36] with sufficiently many iterations.

Figure 6 exhibits the approximate feedback controls and the associated costs from the FIPDE and NNPG methods. From Figure 6 (left), we see that despite the nonconvexity of the control problem, the FIPDE method converges exponentially in the value function approximation as the iteration index m tends to infinity, with the absolute error $\mathcal{O}(0.5^m)$. Compared with the IPDE method, the momentum step in Algorithm 2 accelerates the algorithm’s convergence also in the present nonconvex setting. We refer the reader to Figure 12 in Appendix A.2 for the flocking behaviour of the controlled velocity process obtained by the FIPDE method.

More importantly, the FIPDE method generates a more interpretable approximate feedback control compared to the NNPG method, which helps us understand the mechanism of the optimal control process. Note that for the uncontrolled C-S model (5.5) with $\beta > 0$, the interaction kernel κ indicates that a particle’s velocity is largely affected by the velocities of particles in the nearest neighborhood. Consequently, for a given particle whose velocity is above the average velocity, the further it is away from the population with small velocity, the less internal attraction exists for the particle’s velocity to the average velocity,

and the stronger external intervention is required to induce consensus (and similarly for particle whose velocity is below the average velocity). Such a nonlinear dependence of the optimal feedback control on the variable x is correctly captured by the approximate feedback control ϕ of the FIPDE method, whose values at $t = 0$ are depicted in Figure 6 (middle). Recall that at $t = 0$, the average velocity is 1.5, and particles with velocity below and above 1.5 cluster around the points $x = 1.8$ and $x = 1.2$, respectively. Hence, the absolute magnitude of the function $x \mapsto \phi(0, x, v)$ is minimized near 1.8 if $v > 1.5$ and near 1.2 if $v < 1.5$, which confirms our theoretical understanding of the model. By contrast, the NNPG method produces a less interpretable approximate feedback control with no clear dependence on the variable x , as shown in Figure 6 (right).




Figure 6: Numerical results of the FIPDE and NNPG methods for two-dimensional C-S model with $\beta = 10$; from left to right: convergence of the FIPDE scheme in terms of NAG iterations, the feedback control of the FIPDE method at $t = 0$, and the feedback control of the NNPG method at $t = 0$.

5.2.2 Sparse consensus control of six-dimensional C-S models

This section examines the performance of the FIPDE method (i.e., Algorithm 2) for solving (5.5)-(5.6) with possibly nonsmooth costs in a six-dimensional setting, where the corresponding PDE systems (5.7) are solved by using the residual approximation approach introduced in Section 4. We carry out numerical experiments with $d = 3$, $T = K = 1$, $\sigma = 0.1\mathbb{I}_3$, $\gamma_1 = 0.3$, $(x_0, v_0) \sim \mathcal{U}([0, 1]^6)$ (the uniform distribution on $[0, 1]^6$) but different choices of $\beta, \gamma_2 \geq 0$. For the sake of presentation, we shall only specify the neural network architectures used in our computation, and refer the reader to Appendix A.2 for the detailed implementation of the empirical risk minimization problem (4.5).

Let us start with the LQ MFC problem by choosing $\beta = \gamma_2 = 0$. We initialize Algorithm 2 with $\phi^0 = 0$ and $\tau = 1/5$, and adopt feedforward neural networks with the same architecture to approximate each component of $\phi^m, \psi^m : [0, T] \times \mathbb{R}^6 \rightarrow \mathbb{R}^3$ and $u^m : [0, T] \times \mathbb{R}^6 \rightarrow \mathbb{R}^6$ for all NAG iterations. More precisely, for each m , we have $\phi^m = (\phi_1^\theta, \dots, \phi_3^\theta)$, $\psi^m = (\psi_1^\theta, \dots, \psi_3^\theta)$ and $u^m = (u_1^\theta, \dots, u_6^\theta)$, where $\phi_i^\theta, \psi_i^\theta, u_i^\theta$ are some fully-connected networks with the sigmoid activation function, depth 2 (1 hidden layer), and the dimensions of the input, output and hidden layers being 7, 1 and 20, respectively. We then apply the residual approximation method to solve (5.7) on the spatial domain $\mathcal{D} = [-1, 2]^6$. Note that approximating each component of the solutions individually allows us to capture the heterogeneity among components with shallow networks, and consequently makes the empirical residuals relatively easy to optimize. Moreover, we set $\bar{u} = 0$ on $[0, T] \times \mathcal{D}$ for the decomposition (4.2), since in the present LQ case, there is no obvious candidate solution to (5.7), except the exact solution based on Riccati equations.

We carry out the FIPDE method with 15 NAG iterations and summarize the numerical results in Table 1. It is clear that the approximate controls and their associated costs of the FIPDE method converge rapidly to the optimal control process and the optimal cost function obtained by Riccati equations. Moreover, a sensitivity analysis of the neural network architectures indicates that the convergence of the FIPDE method is very robust with respect to the network depth L , the dimension H of hidden layers, and the (smooth) activation functions. Figure 13 in Appendix A.2 depicts the trajectories of the uncontrolled and controlled velocity processes, which demonstrates that the approximate feedback control of the FIPDE method effectively accelerates the emergence of the flocking behaviour.

We then examine the accuracy of the FIPDE method for approximating the optimal feedback control ϕ^* , which satisfies $\phi^*(t, x, v) = A_t(v - \mathbb{E}[v_t^*])$ for all $(t, x, v) \in [0, T] \times \mathbb{R}^3 \times \mathbb{R}^3$, with $A \in C([0, T]; \mathbb{R}^{3 \times 3})$.

Table 1: Impact of the network architecture on the convergence of the FIPDE method. Shown are: (a) the absolute error $|J(\phi^m) - J^*|$ of the numerical value function $J(\phi^m)$, and (b) the absolute error $(\mathbb{E}[\sup_{t \in [0, T]} |\alpha_t^m - \alpha_t^*|^2])^{\frac{1}{2}}$ of the numerical control process α^m .

	NAG Itr m	1	3	6	9	12	15
$L = 2, H = 20, \varrho = \text{sigm}$	(a)	0.0401	0.0121	0.0050	0.0023	0.0010	0.0007
	(b)	0.2650	0.1549	0.1015	0.0632	0.0400	0.0283
$L = 4, H = 20, \varrho = \text{sigm}$	(a)	0.0396	0.0122	0.0045	0.0019	0.0008	0.0005
	(b)	0.2676	0.1667	0.1082	0.0721	0.0470	0.0346
$L = 2, H = 40, \varrho = \text{sigm}$	(a)	0.0395	0.0127	0.0053	0.0024	0.0012	0.0009
	(b)	0.2638	0.1594	0.1015	0.0663	0.0412	0.0300
$L = 2, H = 20, \varrho = \text{tanh}$	(a)	0.0412	0.0197	0.0086	0.0041	0.0022	0.0013
	(b)	0.2659	0.1936	0.1261	0.0860	0.0557	0.0346

being the solution to a differential Riccati equation, and v^* being the optimal velocity process. As shown in Figure 7 (left), although the FIPDE method is implemented with feedforward neural networks taking inputs (t, x, v) , the resulting approximate feedback control correctly captures the important structures of the optimal control ϕ^* , i.e., the affineness in v and the independence of x . To further highlight the advantage of the PDE-based solver over pure data-driven approaches, we implement the NNPG method with the same neural network architecture (see Appendix A.2 for more details), and compare the resulting feedback control against the FIPDE method and the analytic solution. Comparing Figure 7 (middle) and (right), we see the FIPDE method recovers the exact feedback control accurately on (cross-sections of) the entire space-time domain, while the NNPG method only approximates the exact control on certain sub-domains (depending on the trajectories of the optimal state process), and fails to capture the time dependence of the exact control.




Figure 7: Feedback controls for six-dimensional C-S model with $\beta = \gamma_2 = 0$; from left to right: feedback control of the FIPDE scheme at $t = 0$ and different (x_3, v_3) , feedback controls of the FIPDE scheme and the analytic solution at different (t, v_3) , and feedback controls of the NNPG scheme and the analytic solution at different (t, v_3) , where the remaining components are set to 0.5.

Now we proceed to apply the FIPDE method to the nonsmooth MFC problem (5.5)-(5.6) with $\beta = 0$ and $\gamma_2 \in \{0.1, 0.3\}$. As initial feedback control ϕ^0 for Algorithm 2 we use the approximate feedback control for the LQ problem (with $\gamma_2 = 0$) obtained by the FIPDE method, and choose the stepsize $\tau = 1/5$. At the m -th NAG iteration, the residual approximation (RA) approach in Section 4 is implemented to solve (5.7) on the domain $\mathcal{D} = [-1, 2]^6$. In particular, we decompose the numerical solution into the form $\bar{u} + \tilde{u}^m$, where $\bar{u} : [0, T] \times \mathbb{R}^6 \rightarrow \mathbb{R}^6$ is the approximate decoupling field of the above LQ problem obtained by the FIPDE method, and $\tilde{u}^m : [0, T] \times \mathbb{R}^6 \rightarrow \mathbb{R}^6$ is an unknown nonlinear residual correction. We then approximate the components of \tilde{u}^m by neural networks with the sigmoid activation function, depth 4 (3 hidden layers), and the dimensions of the input, output and hidden layers being 7, 1 and 40, respectively, and determine the optimal neural network representation by minimizing the empirical loss (4.5) with 100 SGD iterations (see Appendix A.2). In the following, we refer to the method as “RA with 100 SGD”.

To demonstrate the efficiency of the RA approach, we compare the performance of the above FIPDE method to those of the Direct Method and the NNPG method. In the Direct Method, we do not decompose numerical solutions in a separable form (i.e., we set $\bar{u} = 0$ in (4.2)), and directly minimize the residuals of (5.7) over neural networks. For each NAG iteration, we choose the same trial functions (4-layer networks with hidden width 40), numbers of training samples and learning rates of the SGD algorithm as those of the FIPDE method, and perform q SGD iterations with different choices of $q \in \mathbb{N}$, which will be referred to as “DM with q SGD” in the following discussion. In the NNPG method, we minimize the objective (5.6) over all feedback controls parametrized by 4-layer networks with the same architecture (see Appendix A.2 for more details).

Figure 8 depicts the convergence of value functions from “RA with 100 SGD” and “DM with q SGD” (with $q \in \{200, 500, 1000\}$) for solving (5.5)-(5.6) with $\gamma_2 = 0.1$, where we take the value function from the NNPG method as a reference value. It clearly demonstrates that the residual approximation requires significantly less number of SGD iterations than the Direct Method for solving the PDE system (5.7) at each NAG iteration; in particular, the approximate feedback controls of “RA with 100 SGD” (the red line) at all NAG iterations consistently achieve lower costs than those of “DM with 1000 SGD” (the blue line). Observe that for each NAG iteration, the computation times of the residual approximation approach and the Direct Method scale linearly with the number of SGD iterations. Thus the residual decomposition (4.2) reduces the total computational time of the Direct Method by roughly a factor of 10. A similar efficiency enhancement of the residual approximation approach has also been observed in the nonsmooth problem (5.5)-(5.6) with $\beta = 0$ and $\gamma_2 = 0.3$.




Figure 8: Efficiency improvement of the residual approximation approach over the Direct Method for six-dimensional C-S model with $\beta = 0$ and $\gamma_2 = 0.1$.

Figure 9 presents the approximate feedback controls of the nonsmooth problem (5.5)-(5.6) with different $\gamma_2 > 0$, obtained by using the FIPDE and NNPG methods. Due to the ℓ_1 -norm in (5.6), the optimal feedback control is nonlinear in v and admits a sparse structure. That is, the optimal decision is to act only on the particles whose velocities are far from the mean and to steer them to consensus, without intervening with the particles near the consensus manifold [19]. As shown in Figure 9 (left) and (middle), the approximate controls of the FIPDE method correctly capture the sparse features of the optimal control, where the zero-control region shrinks as the terminal time approaches and expands as the parameter γ_2 increases. On the other hand, the NNPG method results in a (suboptimal) non-sparse linear strategy, which suggests to control all agents with mild strength. It also fails to capture the time dependence of the optimal control, as already observed in the LQ setting (see Figure 7). Consequently, for $\gamma_2 = 0.3$, the control of the NNPG method has a 2.4% higher expected cost (5.6) than that of the FIPDE method (in terms of the relative error); the expected costs for the NNPG and FIPDE methods are 0.1635 and 0.1596, respectively.

Finally, we apply the FIPDE method to solve the nonconvex MFC problem (5.5)-(5.6) with $\beta = 1$ and $\gamma_2 = 0$. As in the two-dimensional setting studied in Section 5.2.1, the uncontrolled velocity process does not form a consensus, and the optimal feedback control in general depends nonlinearly on (t, x, v) . Similar to the above nonsmooth problem, we use the LQ feedback control as the initial guess ϕ^0 in Algorithm 2, choose the stepsize $\tau = 1/5$, and employ the residual approximation approach to solve (5.7) at all NAG iterations, with the candidate solution \bar{u} being the approximate decoupling field of the LQ problem. The components of the residual \tilde{u}^m and the controls ϕ^m, ψ^m are approximated by 7-layer neural networks




Figure 9: Feedback controls for six-dimensional C-S models with $\beta = 0$ and $\gamma_2 > 0$, where the controls are evaluated at different (t, v_3) with the remaining components being 0.5.

with hidden width 60 (see Appendix A.2 for more details). Figure 10 compares the uncontrolled velocity process and the controlled velocity process obtained by the FIPDE method (with 6 NAG iterations), which clearly demonstrates the effectiveness of the FIPDE method on inducing consensus. Compared with the NNPG method (implemented with the same neural networks), the FIPDE method results in a similar value function but a better consensus at the terminal time; the value functions of the FIPDE and NNPG methods are 0.111 and 0.114, respectively, while the final variances $\mathbb{E}[|v_T - \mathbb{E}[v_T]|^2]$ obtained by the FIPDE and NNPG methods are 0.007 and 0.01, respectively. Moreover, similar to the above examples in Figures 7 and 9, the FIPDE method captures the nonlinear time dependence of the optimal control, while the NNPG method leads to an approximate control that is constant in t (plots omitted).




Figure 10: Uncontrolled and controlled six-dimensional C-S models with $\beta = 1$ at different times.

Appendix A Supplementary materials for Section 5

A.1 Supplementary materials for Section 5.1

Convergence of the NAG iteration. Figure 11 presents the expected costs of the approximate feedback controls obtained by all NAG iterations of the FIPDE and EMReg methods for the optimal liquidation problem (5.1) with $Q_0 \in \mathcal{U}(1, 2)$, $k_1 = 1$ and $k_2 \in \{0, 1\}$. To estimate the optimal expected cost for the LQ setting (with $k_2 = 0$), we implement the exact feedback control (5.4) with time stepsize $\Delta t = 1/50$, and replace the expectation in (5.1) by the empirical average over 10^6 sample trajectories of the state processes. It shows that for both the LQ and nonsmooth MFC problems, the FIPDE and EMReg methods give convergent approximations to the optimal cost functions as the number of NAG iterations tends to infinity. Moreover, a linear regression of the data indicates that the FIPDE method approximates the value function with an absolute error of the magnitude $\mathcal{O}(m^{-2.7})$ for the LQ case.

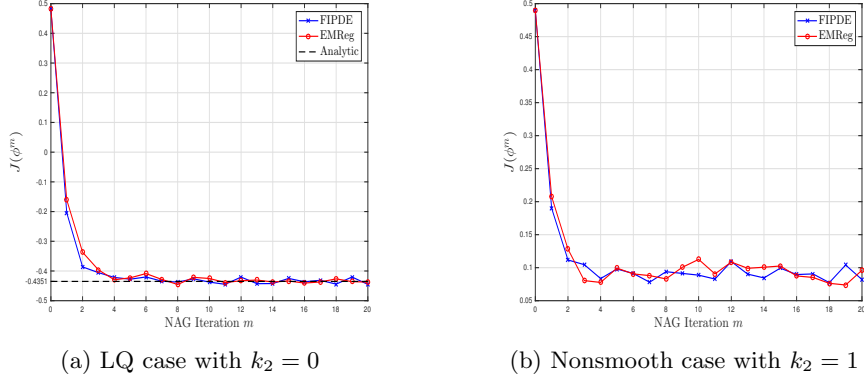


Figure 11: Convergence of the FIPDE and EMReg methods for approximating the value functions of optimal liquidation problems (5.1) in terms of NAG iteration.

Computational time. The FIPDE method is implemented by using MATLAB R2016b on a laptop with 2.2GHz 4-core Intel Core i7 processor and 16 GB memory. The computation takes around 1 second per NAG iteration. The EMReg method is implemented by using MATLAB R2020b on a PC with 2.1GHz 6-core Intel Core i5 processor and 16 GB memory. The computation takes around 5 minutes per NAG iteration based on our implementation.

A.2 Supplementary materials for Section 5.2

Effectiveness of the FIPDE method for two-dimensional models. Figure 12 compares the uncontrolled two-dimensional C-S model with $\beta = 10$ and the controlled model obtained by the FIPDE method, where the scatter plots are generated based on 10^4 simulated trajectories. One can clearly observe that the uncontrolled velocity process does not admit a time-asymptotic flocking behaviour, and the feedback control from the FIPDE method effectively induces the consensus of the velocity process.

Effectiveness of the FIPDE method for six-dimensional models with $\beta = 0$. Figure 12 compares the uncontrolled six-dimensional C-S model with $\beta = 0$ and the controlled model obtained by the FIPDE method with quadratic costs (i.e., $\gamma_2 = 0$ in (5.6)), where the scatter plots are generated based on 5×10^3 simulated trajectories. One can clearly observe that the feedback control from the FIPDE method accelerates the consensus of the velocity process.

Implementation of the NNPG method for two-dimensional models. We implement the neural network-based policy gradient (NNPG) method using PyTorch for the optimal control of two-dimensional C-S models with $\beta = 10$ as follows. The set of trial functions for feedback controls consists of 7 layer fully-connected networks with the sigmoid activation function and the dimensions of input, hidden and output layers being equal to 3, 70 and 1, respectively. The optimal weights θ of the network are obtained by running the Adam algorithm [36] with 5×10^4 iterations and a decaying learning rate scheduler. At each iteration, we generate 10^2 trajectories of (5.5) by using (2.24) with time stepsize $\Delta t = 0.02$, approximate the functional J in (5.6) by the empirical average of simulated trajectories, and perform gradient descent based on the empirical loss. The learning rate is initialized at 0.0005 and will be reduced by a factor of 0.3 if the loss does not decrease after 10 iterations. The minimum learning rate is set to be 10^{-6} .

Implementation of the NNPG method for six-dimensional models. We state the PyTorch implementation of the NNPG method for six-dimensional C-S models with different choices of β, γ_2 used in Section 5.2. As the chosen neural network architecture for each example has been specified in Section 5.2, it remains to discuss the implementation of the Adam algorithm. For $\beta = 0$ and $\gamma_2 \in \{0, 0.1, 0.3\}$, at each SGD iteration, we generate 10^3 trajectories of (5.5) by using (2.24) with time stepsize $\Delta t = 0.02$, and perform gradient descent based on an empirical approximation of J in (5.6). We choose the same “reduce on plateau” learning rate scheduler as the above two-dimensional case and perform a sufficiently large number of iterations until the empirical loss stabilizes. For $\beta = 1$ and $\gamma_2 = 0$, we reduce the

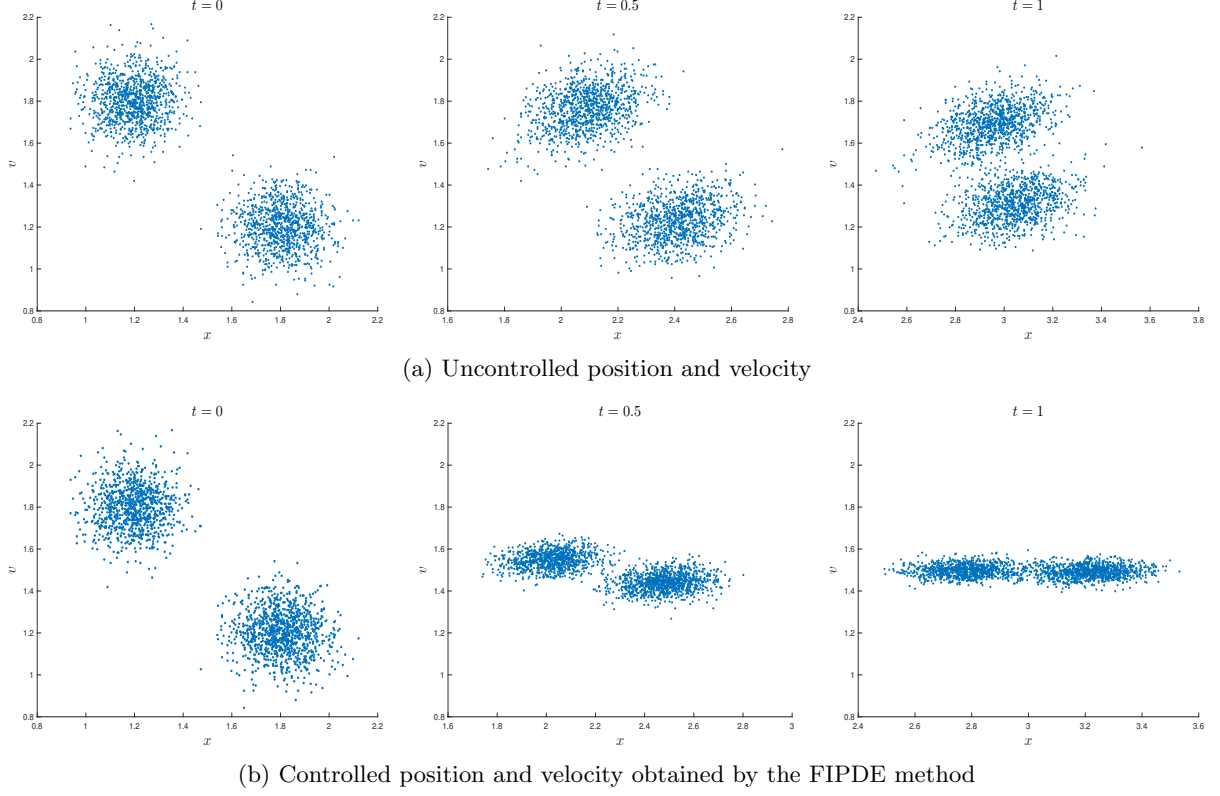


Figure 12: Uncontrolled and controlled two-dimensional C-S models with $\beta = 10$ at different times.

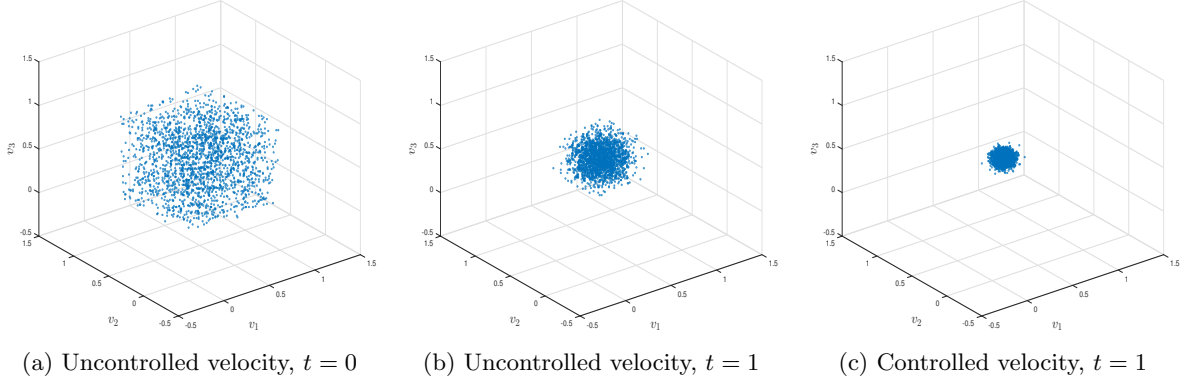


Figure 13: Uncontrolled and controlled six-dimensional C-S models with $\beta = 0$ at different times.

number of simulated trajectories per iteration to 10^2 and keep the remaining configurations. This helps to accommodate the increasing computational cost of simulating (5.5) with $\beta > 0$.

Implementation of the FIPDE method for six-dimensional models. We state the PyTorch implementation of the FIPDE method for six-dimensional C-S models with different choices of β, γ_2 used in Section 5.2. Before stating the configuration details, we remark that the algorithm’s hyperparameters have not been optimally tuned and hence the following choices may not be the optimal combination. The neural network architectures for all examples have been specified in Section 5.2.

For all NAG iterations, we sample N trajectories of (5.5) by using (2.24) with $\Delta t = 0.01$, where $N = 4 \times 10^4$ for $\beta = \gamma_2 = 0$, and $N = 5 \times 10^3$ for the remaining cases. Then, for all choices of β, γ_2 , we apply the Adam algorithm to solve the corresponding risk minimization problem (4.4) with the initial learning rate 0.005. The learning rate will be decreased by employing the same “reduce on plateau” scheduler as the above NNPG method for the examples with $\beta = 0, \gamma_2 \in \{0, 0.1, 0.3\}$, while for the

example with $\beta = 1, \gamma_2 = 0$ it will be decreased by a factor of 0.8 for every 5000 SGD iterations.

At each SGD iteration, a mini-batch of points with size 20 is drawn by following uniform distributions from the interior and boundaries of the domain (i.e., $N_{\text{in}} = N_{\text{ter}} = N_{\text{bdy}} = 20$ in (4.5)). Based on these samples, we compute the empirical loss (4.5) with certain $\eta_1, \eta_2 > 0$ depending on the model parameters. In particular, for $\beta = \gamma_2 = 0$, we choose $\eta_1 = 5, \eta_2 = 10$ for the first 6 NAG iterations and $\eta_1 = 2, \eta_2 = 5$ for the remaining NAG iterations to balance the interior and boundary losses, while for $\beta = 0, \gamma_2 = \{0, 1, 0.3\}$ and $\beta = 1, \gamma_2 = 0$, we choose $\eta_1 = 1, \eta_2 = 1/400$ and $\eta_1 = 1, \eta_2 = 1/800$, respectively, across all NAG iterations. The total number of Adam iterations is chosen as 9000 for $\beta = \gamma_2 = 0$, 100 for $\beta = 0, \gamma_2 \in \{0.1, 0.3\}$, and 20000 for $\beta = 1, \gamma_2 = 0$.

To solve the supervised learning problem (4.6), we carry out the Adam algorithm with initial learning rate 0.001, which will be reduced by employing the same scheduler as that for solving (4.4). We shall randomly draw a mini-batch of points with size 20 at each iteration, and run the algorithm with sufficiently many iterations until the loss stabilizes.

Computational time. For the two-dimensional C-S models, the FIPDE method is implemented in MATLAB R2016b on a laptop with 2.2GHz 4-core Intel Core i7 processor and 16 GB memory. The computation with 15 NAG iterations takes around 30 seconds for $\beta = 0$ and around 10 minutes for $\beta = 10$, due to an increased computational cost in evaluating the interaction kernel κ for general $\beta > 0$. The NNPG method is implemented in Python 3.8.5 on a PC with 2.1GHz 6-core Intel Core i5 processor and 16 GB memory. The computation with 5×10^4 Adam iterations takes around 2 hours.

For the six-dimensional C-S models, both the FIPDE and NNPG methods are implemented by using Python 3.8.5 on a PC with 2.1GHz 6-core Intel Core i5 processor and 16 GB memory. The NNPG method typically takes around 20-40 minutes, where the precise computation time depends on the complexity of the problem and the number of SGD iterations used in the simulation. The FIPDE method takes around 20 minutes per NAG iteration for $\beta = 0$, and around 8 hours per NAG iteration for $\beta = 1$. Note that these times can be shortened if we fine tune the hyperparameters (e.g., the number of SGD iterations) and perform the computation on GPUs.

Acknowledgements

Wolfgang Stockinger is supported by a special Upper Austrian Government grant.

References

- [1] B. ACCIAIO, J. BACKHOFF-VERAGUAS, AND R. CARMONA, *Extended mean field control problems: stochastic maximum principle and transport perspective*, SIAM Journal on Control and Optimization, 57 (2019), pp. 3666–3693.
- [2] Y. ACHDOU AND M. LAURIÈRE, *On the system of partial differential equations arising in mean field type control*, arXiv preprint arXiv:1503.05044, (2015).
- [3] ———, *Mean field type control with congestion (II): An augmented Lagrangian method*, Applied Mathematics & Optimization, 74 (2016), pp. 535–578.
- [4] ———, *Mean field games and applications: Numerical aspects*, Mean Field Games, (2020), pp. 249–307.
- [5] G. ALBI, M. HERTY, D. KALISE, AND C. SEGALA, *Moment-driven predictive control of mean-field collective dynamics*, arXiv preprint arXiv:2101.01970, (2021).
- [6] A. ANGIULI, C. V. GRAVES, H. LI, J.-F. CHASSAGNEUX, F. DELARUE, AND R. CARMONA, *Cemracs 2017: numerical probabilistic approach to MFG*, ESAIM: Proceedings and Surveys, 65 (2019), pp. 84–113.
- [7] R. ARCHIBALD, F. BAO, J. YONG, AND T. ZHOU, *An efficient numerical algorithm for solving data driven feedback control problems*, Journal of Scientific Computing, 85 (2020), pp. 1–27.

- [8] B. AZMI, D. KALISE, AND K. KUNISCH, *Optimal feedback law recovery by gradient-augmented sparse polynomial regression*, Journal of Machine Learning Research, 22 (2021), pp. 1–32.
- [9] R. BAIO, M. BONGINI, J. A. CARRILLO, AND D. KALISE, *Optimal consensus control of the Cucker-Smale model*, IFAC-PapersOnLine, 51 (2018), pp. 1–6.
- [10] J. BAO, C. REISINGER, P. REN, AND W. STOCKINGER, *First-order convergence of Milstein schemes for McKean–Vlasov equations and interacting particle systems*, Proceedings of the Royal Society A, 477 (2021), p. 20200258.
- [11] G. BARLES AND P. E. SOUGANIDIS, *Convergence of approximation schemes for fully nonlinear second order equations*, Asymptotic Analysis, 4 (1991), pp. 271–283.
- [12] M. BASEI AND H. PHAM, *Linear-quadratic McKean–Vlasov stochastic control problems with random coefficients on finite and infinite horizon, and applications*, arXiv preprint arXiv:1711.09390, (2017).
- [13] A. BECK, *First-order methods in optimization*, SIAM, 2017.
- [14] A. BECK AND M. TEOULLE, *Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems*, IEEE transactions on image processing, 18 (2009), pp. 2419–2434.
- [15] ———, *A fast iterative shrinkage-thresholding algorithm for linear inverse problems*, SIAM Journal on Imaging Sciences, 2 (2009), pp. 183–202.
- [16] C. BENDER AND J. ZHANG, *Time discretization and Markovian iteration for coupled FBSDEs*, The Annals of Applied Probability, 18 (2008), pp. 143–177.
- [17] I. H. BISWAS, E. R. JAKOBSEN, AND K. H. KARLSEN, *Difference-quadrature schemes for nonlinear degenerate parabolic integro-PDE*, SIAM Journal on Numerical Analysis, 48 (2010), pp. 1110–1135.
- [18] M. BOSSY AND D. TALAY, *A stochastic particle method for the McKean–Vlasov and the Burgers equation*, Mathematics of Computation, 66 (1997), pp. 157–192.
- [19] M. CAPONIGRO, M. FORNASIER, B. PICCOLI, AND E. TRÉLAT, *Sparse stabilization and optimal control of the Cucker-Smale model*, Mathematical Control and Related Fields, 3 (2013), pp. 447–466.
- [20] R. CARMONA AND F. DELARUE, *Probabilistic Theory of Mean Field Games with Applications I*, Springer, 2018.
- [21] R. CARMONA AND M. LAURIÈRE, *Convergence analysis of machine learning algorithms for the numerical solution of mean field control and games: II—the finite horizon case*, arXiv preprint arXiv:1908.01613, (2019).
- [22] ———, *Convergence analysis of machine learning algorithms for the numerical solution of mean field control and games I:—the ergodic case*, SIAM Journal on Numerical Analysis, 59 (2021), pp. 1455–1485.
- [23] J. A. CARRILLO, M. FORNASIER, J. ROSADO, AND G. TOSCANI, *Asymptotic flocking dynamics for the kinetic Cucker-Smale model*, SIAM Journal on Mathematical Analysis, 42 (2010), pp. 218–236.
- [24] F. CUCKER AND S. SMALE, *Emergent behavior in flocks*, IEEE Transactions on Automatic Control, 52 (2007), pp. 852–862.
- [25] K. DEBRABANT AND E. JAKOBSEN, *Semi-lagrangian schemes for linear and fully non-linear diffusion equations*, Mathematics of Computation, 82 (2013), pp. 1433–1462.
- [26] W. E, J. HAN, AND A. JENTZEN, *Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations*, Communications in Mathematics and Statistics, 5 (2017), pp. 349–380.
- [27] M. GERMAIN, M. LAURIÈRE, H. PHAM, AND X. WARIN, *Deepsets and their derivative networks for solving symmetric PDEs*, arXiv preprint arXiv:2103.00838, (2021).

- [28] M. GERMAIN, J. MIKAEL, AND X. WARIN, *Numerical resolution of McKean–Vlasov FBSDEs using neural networks*, arXiv preprint arXiv:1909.12678, (2019).
- [29] S. GHADIMI AND G. LAN, *Accelerated gradient methods for nonconvex nonlinear and stochastic programming*, Mathematical Programming, 156 (2016), pp. 59–99.
- [30] M. GIEGRICH, C. REISINGER, AND Y. ZHANG, *Convergence of policy gradient methods for finite-horizon stochastic linear-quadratic control problems*, arXiv preprint arXiv:2211.00617, (2022).
- [31] E. GOBET AND M. GRANGEREAU, *Extended McKean–Vlasov optimal stochastic control applied to smart grid management*, (2019).
- [32] O. GUÉANT, *The Financial Mathematics of Market Liquidity: From optimal execution to market making*, vol. 33, CRC Press, 2016.
- [33] X. GUO, A. HU, AND Y. ZHANG, *Reinforcement learning for linear-convex models with jumps via stability analysis of feedback controls*, arXiv preprint arXiv:2104.09311, (2021).
- [34] K. ITO, C. REISINGER, AND Y. ZHANG, *A neural network-based policy iteration algorithm with global H^2 -superlinear convergence for stochastic games on domains*, Foundations of Computational Mathematics, 21 (2021), pp. 331–374.
- [35] B. KERIMKULOV, D. ŠIŠKA, AND L. SZPRUCH, *A modified MSA for stochastic control problems*, Applied Mathematics & Optimization, (2021), pp. 1–20.
- [36] D. P. KINGMA AND J. BA, *Adam: A method for stochastic optimization*, arXiv preprint arXiv:1412.6980, (2014).
- [37] H. J. K. KUSHNER, H. J. KUSHNER, P. G. DUPUIS, AND P. DUPUIS, *Numerical methods for stochastic control problems in continuous time*, vol. 24, Springer Science & Business Media, 2001.
- [38] J.-P. LEMOR, E. GOBET, AND X. WARIN, *Rate of convergence of an empirical regression method for solving generalized backward stochastic differential equations*, Bernoulli, 12 (2006), pp. 889–916.
- [39] H. LI AND Z. LIN, *Accelerated proximal gradient methods for nonconvex programming*, Advances in neural information processing systems, 28 (2015).
- [40] Q. LI, Y. ZHOU, Y. LIANG, AND P. K. VARSHNEY, *Convergence analysis of proximal gradient with momentum for nonconvex optimization*, in International Conference on Machine Learning, PMLR, 2017, pp. 2111–2119.
- [41] T. MILLER, *Explanation in artificial intelligence: Insights from the social sciences*, Artificial Intelligence, 267 (2019), pp. 1–38.
- [42] Y. NESTEROV, *Introductory lectures on convex optimization: A basic course*, vol. 87, Springer Science & Business Media, 2003.
- [43] M. NOURIAN, P. E. CAINES, AND R. P. MALHAMÉ, *Mean field analysis of controlled Cucker–Smale type flocking: Linear analysis and perturbation equations*, IFAC Proceedings Volumes, 44 (2011), pp. 4471–4476.
- [44] É. PARDOUX, *Backward stochastic differential equations and viscosity solutions of systems of semi-linear parabolic and elliptic PDEs of second order*, in Stochastic Analysis and Related Topics VI, Springer, 1998, pp. 79–127.
- [45] L. PFEIFFER, *Numerical methods for mean-field-type optimal control problems*, arXiv preprint arXiv:1703.10001, (2017).
- [46] H. PHAM AND X. WEI, *Bellman equation and viscosity solutions for mean-field stochastic control problem*, ESAIM: Control, Optimisation and Calculus of Variations, 24 (2018), pp. 437–461.
- [47] C. REISINGER, W. STOCKINGER, AND Y. ZHANG, *Optimal regularity of extended mean field controls and their piecewise constant approximation*, arXiv preprint arXiv:2009.08175v2, (2020).

- [48] ———, *A posteriori error estimates for fully coupled McKean–Vlasov forward-backward SDEs*, arXiv preprint arXiv:2007.07731, (2020).
- [49] ———, *Linear convergence of a policy gradient method for finite horizon continuous time stochastic control problems*, arXiv preprint arXiv:2203.11758, (2022).
- [50] C. REISINGER AND Y. ZHANG, *Rectified deep neural networks overcome the curse of dimensionality for nonsmooth value functions in zero-sum games of nonlinear stiff systems*, Analysis and Applications, 18 (2020), pp. 951–999.
- [51] M. SCHMIDT, N. ROUX, AND F. BACH, *Convergence rates of inexact proximal-gradient methods for convex optimization*, Advances in neural information processing systems, 24 (2011).
- [52] J. SIRIGNANO AND K. SPILIOPOULOS, *DGM: A deep learning algorithm for solving partial differential equations*, Journal of Computational Physics, 375 (2018), pp. 1339–1364.
- [53] D. ŠIŠKA AND Ł. SZPRUCH, *Gradient flows for regularized stochastic control problems*, arXiv preprint arXiv:2006.05956, (2020).
- [54] J. YONG, *Linear-quadratic optimal control problems for mean-field stochastic differential equations*, SIAM Journal on Control and Optimization, 51 (2013), pp. 2809–2838.