



Review of Mean-Field Theory for the Robotic Swarms Common Target Problem

Uma Revisão da Teoria do Campo Médio para o Problema do Alvo Comum em Enxames Robóticos

Yuri Tavares dos Passos¹ João Victor Neves de Souza Nunes²

Received: August 15, 2025

Received in revised form: October 24, 2025

Accepted: October 30, 2025

Available online: November 30, 2025

ABSTRACT

This review provides a literature overview of the usage of Mean-Field Theory (MFT) for modelling the performance estimation of algorithms solving the common target problem in robotic swarm navigation. This problem involves coordinating multiple robots to converge simultaneously towards a single target, which can result in congestion and degraded performance. Its main objective is to minimise the maximum time required for all robots to reach and then depart from the target. The challenge lies in ensuring effective coordination among the robots to prevent conflicts and optimise resource utilisation. MFT may accurately predict navigation metrics, such as agent arrival and departure times, target location density, and system throughput under congestion, and it may enable control strategies based on distributed feedback. This literature overview illustrates that MFT is an effective and promising mathematical tool to mitigate the issues of high dimensionality and complexity of interactions present in systems consisting of a large number of agents.

keywords robotic swarms, mean-field theory, common target problem, differential equations

RESUMO

Esta revisão fornece uma visão geral da literatura sobre o uso da Teoria de Campo Médio (Mean-Field Theory — MFT) para modelar as estimativas de desempenho de algoritmos que resolvem o problema do alvo comum na navegação de enxames robóticos. Este problema envolve a coordenação de múltiplos robôs convergindo simultaneamente em direção a um único alvo, o que pode resultar em congestionamento e desempenho degradado. Seu principal objetivo é minimizar o tempo máximo necessário para que todos os robôs cheguem ao alvo e retornem. O desafio está em garantir uma coordenação eficaz entre os robôs para evitar conflitos e otimizar a utilização de recursos. MFT pode prever com precisão métricas de navegação, como tempos de chegada e partida dos agentes, densidade de localização do alvo e rendimento do sistema sob congestionamento e possibilitar estratégias de controle baseadas na utilização de feedback distribuído. A visão geral da literatura ilustra que a MFT fornece uma ferramenta matemática eficaz e promissora para mitigar os problemas de alta dimensionalidade e complexidade das interações presentes em sistemas compostos por um grande número de agentes.

palavras-chave enxames de robôs, teoria de campo médio, problema do alvo comum, equações diferenciais

¹Prof. Dr., Center for Exact and Technological Sciences, UFRB, Cruz das Almas, BA, Brazil. yuri.passos@ufrb.edu.br

²Graduate student, Center for Exact and Technological Sciences, UFRB, Cruz das Almas, BA, Brazil. tiokarin@aluno.ufrb.edu.br

Introduction

Robotics is an engineering and computer science field dedicated to the design and production of robots. The aim of robotics is to design machines that can carry out certain tasks and interact with the world, augmenting human abilities and systematising processes.

The domain of robotics spans a broad spectrum of fields, ranging from industrial robots to medical and space robots. The use of swarm-based approaches in the development of robotics systems has attracted long-standing interest (Bonabeau et al., 1999; Couzin et al., 2002). Swarms consists of large numbers of collaborating robots to conduct complicated tasks. One of the major challenges facing this field is the formulation of modelling and control methods that scale favourably with swarm size. Considering the swarm as a continuum presents one solution to this problem.

Among the complex tasks addressed in robotic swarms, the common target problem stands out as a particularly challenging scenario that directly exemplifies the need for scalable modelling and control. A pervasive and highly challenging issue in robotic swarm navigation under such conditions is the *Common Target Problem*. This problem arises when multiple robots attempt to reach the same target simultaneously, creating congestion scenarios that may degrade performance (Passos et al., 2023).

The core challenge is coordinating the convergence of large groups of robots into a confined space; any lack of coordination would lead to conflicts and inefficient utilisation of resources. Although the mathematical formulation and constraint definition are elucidated, optimising the dynamics of long-term experiments such as the common target problem poses substantial control difficulties. The escalating complexity associated with an increasing number of robots frequently renders centralised control strategies impractical, thereby underscoring the necessity for approaches that enhance system scalability and robustness.

Therefore, this study focuses on this specific problem within robotic swarms. Let a swarm of N robots be modelled such that the dynamics of each robot $i \in \{1, \dots, N\}$ at position $\mathbf{r}_i(t) = (r_{ix}(t), r_{iy}(t), r_{iz}(t))$ at time t are governed by attractive and repulsive forces. Let $\dot{\mathbf{r}}_i = \frac{d\mathbf{r}_i}{dt}$ and $\ddot{\mathbf{r}}_i = \frac{d\dot{\mathbf{r}}_i}{dt}$. The force $\mathbf{F}_A(\mathbf{T})$ attracts the robot towards a target located at \mathbf{T} while \mathbf{F}_R is the repulsive force applied to avoid bumping into other robots. The working and target areas are circles with radius D and ϵ , respectively, centred on \mathbf{T} with $D > \epsilon$. Let t_{ti} and t_{fi} be the time to reach the target area and leave the working region, respectively, and $t_i = t_{ti} + t_{fi}$. After a robot leaves the working area, it proceeds to a new target located at \mathbf{N}_i . All robots are assumed to be circles of radius d_r . Hence, they have to stay apart from each other by at least $2d_r$. The maximum time between all robots, $\max_{i \in \{1, \dots, N\}} t_i$, and their total work must be minimised, that is,

$$\begin{aligned} \text{minimise } J = \max_{i \in \{1, \dots, N\}} (t_{ti} + t_{fi}) + \sum_{i=1}^N \int_{t_0}^{t_{ti}} \phi_1 \cdot \dot{\mathbf{r}}_i dt + \int_{t_{ti}}^{t_{fi}} \phi_2 \cdot \dot{\mathbf{r}}_i dt \text{ with} \\ \phi_1 = K_{res} \frac{\mathbf{F}_A(\mathbf{T}) + \mathbf{F}_R}{\|\mathbf{F}_A(\mathbf{T}) + \mathbf{F}_R\|} - K_{dp} \dot{\mathbf{r}}_i(t) \text{ and } \phi_2 = K_{res} \frac{\mathbf{F}_A(\mathbf{N}_i) + \mathbf{F}_R}{\|\mathbf{F}_A(\mathbf{N}_i) + \mathbf{F}_R\|} - K_{dp} \dot{\mathbf{r}}_i(t), \\ \text{subject to } \forall i \in \{1, \dots, N\} : \|\mathbf{r}_i(t_{ti}) - \mathbf{T}\| = \epsilon, \quad \|\mathbf{r}_i(t_{fi}) - \mathbf{T}\| = D \\ \text{and } \forall j \in \{1, \dots, N\}, j \neq i : \|\mathbf{r}_i - \mathbf{r}_j\| \geq 2d_r. \end{aligned} \quad (1)$$

The coefficient K_{res} determines the magnitude of the resultant force obtained by the attractive and repulsive forces, while K_{dp} governs motion dissipation, reducing oscillations and smoothing trajectories as the robots move towards their targets. In essence, the equation (1) quantifies the overall cost or effort for a robotic swarm to complete a common task, specifically aiming to minimise the longest time any single robot takes to reach and return from a target, while also considering the accumulated cost of movement over time.

Typically, two categories of optimisation criteria can be used in task allocation to achieve a stated system objective: maximisation (reward-based) or minimisation (cost-based). The first category is employed when there is a reward for assigning a task to any robot, while the second is more applicable in practice when a cost exists for performing a task and can be further subcategorised as follows. To tackle these optimal control challenges, a variety of strategies have been proposed for traffic management and collision avoidance within multi-robot environments.

Within this context, Mean-Field Theory (MFT) provides a robust mathematical tool for modelling collective behaviour by aggregating all agents (Bensoussan et al., 2013; Carmona & Delarue, 2018; Lasry & Lions, 2007), replacing intricate, agent-specific dynamics with macroscopic approximations often expressed through partial and ordinary differential equations. Therefore, an analysis of MFT-based approaches for addressing the common target problem in robotic swarms is presented in this work. These observations lead to the objective of this paper, which is guided by two primary questions:

- How has MFT been utilised in the modelling and performance evaluation of algorithms designed for addressing the common target problem in robotic swarms?
- What are the principal challenges, advantages, and potential applications identified in the literature?

Every subsequent step, from the development of a search plan to data analysis, adhered to these questions as guiding principles.

Materials and methods

To guarantee reproducibility, the following eligibility criteria were defined:

1. Studies addressing modelling, control, performance estimation, or behaviour analysis using MFT in multi-agent systems focused on robotic swarms.
2. Only research published in English was included, with no restrictions on publication date to encompass both recent advances and foundational work.


A systematic search strategy was developed to identify relevant studies. Search terms included combinations of “Mean-Field Control”, “Mean-Field Games”, or “Collective Dynamics” with “Robotic Swarms”, “Multi-agent Systems”, or “Collective Robotics”, alongside keywords such as “Performance Estimation”, “Modelling”, “Control Strategies”, “Behaviour Analysis”, “Throughput”, and “Congestion”. To maintain focus, works that did not explicitly apply MFT to robotic swarms were excluded, as were those from unrelated domains, extended abstracts without full texts, and duplicates.

Results and discussion

The principal significance of MFT resides in its capability to circumvent the curse of dimensionality inherent to the control of numerous interacting agents by representing the system as a continuous probability distribution. This abstraction markedly diminishes the computational strain, facilitating the design of scalable and computationally tractable control strategies.

By concentrating upon the average behaviour, MFT enables the prediction and control of aggregate swarm dynamics without the need to track the individual state of every robot. Figure 1 illustrates this behaviour, showing both the agent’s influence on the collective and the simplification that the mean-field formulation introduces into the interaction structure.

Figure 1 - The grid represents a multi-agent system where each node (agent) is influenced by the mean effect of its neighbours (green area), simplifying many agent interactions into two agent problems.



Adapted from "Mean field multi-agent reinforcement learning" by Y. Yang, R. Luo, M. Li, M. Zhou, W. Zhang & J. Wang, 2018, Proceedings of the 35th International Conference on Machine Learning.

MFT treats energy as a shared resource, formulating optimisation problems that minimise total swarm power rather than individual consumption, yielding efficient trajectories, balanced density and power-saving behaviours that extend mission life. The mathematical engines that turn local rules into global behaviour were examined.

MFT elevates performance in complex tasks, for instance, area coverage, search and rescue (Ménec, 2024) and environmental monitoring, by distilling swarm behaviour into concise mathematical forecasts. It supplies unified frameworks that predict collective motion, enabling cooperative strategies for thorough coverage, precise localisation and efficient resource use. Framing tasks as Mean-Field Control or Game problems yields global strategies that raise success rates while curbing congestion and redundant effort (Ornia et al., 2022). MFT further underpins task allocation (Cui et al., 2023), coverage, self-assembly, consensus, flocking (Elamvazhuthi et al., 2021) and strategic interactions.

Continuous settings rely on Fokker–Planck Partial Differential Equations (PDE). When agents pursue selfish optima, a Mean-Field Game emerges, coupling Hamilton–Jacobi–Bellman and Fokker–Planck equations. Essentially, these equations model how the density or distribution of a large population of agents evolves over time and space, accounting for their movement, ensuring that the total number of agents is conserved within the system (Burger et al., 2014).

Thus, MFT moves beyond theory to become indispensable for scalable, energy-aware and task-focused swarms. It replaces individual control with unified continuum models.

MFT becomes an adequate method of simulation and control of collective behaviours in robotic swarms using modelling techniques borrowed from natural sciences—for example, astrophysics (Hazra et al., 2023), biology (Couzin et al., 2002) and statistical mechanics (Vicsek et al., 1995). Thus, MFT assumes that individual agent behaviour can be approximated by the mean behaviour of all agents within a system. Hence, simplifying the macroscopic description and making it much more accessible for analysis. Under this approximation, two major theoretical frameworks have emerged developed: Mean-Field Game (MFG) and Mean-Field Control (MFC) problems.

Despite substantial research growth and notable progress in applying MFT to robotic swarms, a more focused synthesis in the review literature remains valuable. While existing reviews have comprehensively explored MFT in robotics (Elamvazhuthi & Berman, 2019), their scope has not specifically extended to a thorough assessment of MFT's role in performance evaluation for common target scenarios, nor have they systematically compared different control strategies amidst congestion and conflicting agent preferences.

Consequently, although the foundational context of MFT in multi-agent systems is well covered, a dedicated review addressing the practical implications of MFT in handling these complex swarm behaviours, particularly concerning performance evaluation and specific control strategy comparisons, is timely and complementary to the existing body of work.

Following the established framework, the ensuing sections are systematically structured according to the classification of mean-field models. The initial section delves into the application in robotic swarms of two primary categories of mean-field models, infinite-dimensional continuous and finite-dimensional discrete, facilitating a comparative analysis between these two paradigms.

The second section addresses mean-field models that originate from the fields of physics and statistical mechanics, which elucidate the mechanisms by which complex physical concepts are examined via mean-field approximations. Having framed the challenge, how MFT enables the modelling and control of collective behaviours in robotic swarms is analysed below.

This review uniquely bridges the gap between MFT theory and practical swarm robotics challenges, providing validated solutions for congestion and collision avoidance areas largely overlooked in existing literature. To better clarify the distinction between existing reviews on Mean Field Theory applied to robotic swarms and the current review, the main differences can be observed in Table 1.

Mean-field models in robotic swarms —————

Building on this theoretical framework, it is evident that the inherent complexities of robotic swarms, which are characterised by a large number of interacting agents, provide both tremendous opportunities and difficult obstacles. While this paradigm has great real-world application potential, notably in terms of scalability, robustness, and decentralised function, it is also fraught with complications and challenges.

Table 1 - Thematic and Contribution-Focussed Comparison of Literature Reviews

| Reference | Main Objectives | Main Themes | Practical Implications |
|--------------------------------|--|--|---|
| Elamvazhuthi and Berman (2019) | To provide a fundamental understanding and overview of MFT applications in multi-agent systems. | <ul style="list-style-type: none"> Theoretical foundations of MFT. General applicability in robotics. | The authors articulate theoretical advantages; however, they establish a less explicit connection to detailed validation and practical application. |
| Vicsek et al. (1995) | To model collective motion and phase transitions in systems of self-driven particles using statistical physics. | <ul style="list-style-type: none"> Local alignment rules. Phase transition to global order. Noise-driven fluctuations. | Demonstrates how local interaction rules lead to global coherence; forms the basis for swarm coordination models in robotics. |
| Bensoussan et al. (2013) | To introduce mean-field games and mean-field type control theory. | <ul style="list-style-type: none"> Mathematical foundations of MFG and MFC. Applications to large-scale systems. | Provides rigorous theoretical background, but lacks direct application to robotic swarms. |
| Carmona and Delarue (2018) | To present a probabilistic theory of mean-field games and applications. | <ul style="list-style-type: none"> Stochastic analysis of MFG. Coupled HJB and FP equations. | High mathematical complexity; limited focus on robotic systems or empirical validation. |
| Lasry and Lions (2007) | To establish the foundational theory of mean-field games. | <ul style="list-style-type: none"> Nash equilibria in large populations. Coupled PDE systems (HJB + FP). | Introduces core MFG concepts; not tailored to robotic swarm applications. |
| Transporte et al. (2022) | To propose a top-down control strategy for robotic swarms using mean-field PDEs. | <ul style="list-style-type: none"> Density control via velocity fields. Input-to-State Stability (ISS). | Demonstrates scalable control with ISS guarantees; relevant to common target problems. |
| This Review | To provide a pragmatic and concentrated analysis of the utility of MFT in addressing the challenges related to performance estimation and control in robotic swarms. | <ul style="list-style-type: none"> Performance assessment using MFT. Comparison of control strategies. Prediction of navigation metrics. Validation via direct simulation. | The current review serves as a conduit between the overarching principles of MFT and its empirically validated application in addressing particular congestion and conflict issues within robotic swarms, a domain that has not been comprehensively analysed in preceding reviews. |

Efficiently managing the synchronisation of large, distributed teams is crucial for maximising the potential of these systems. This fundamental duality, the appealing promise and the complex coordination challenges, serves as a foundation for recognising the need of advanced modelling and control frameworks in this rapidly evolving domain.

The transition from individual dynamics to collective behaviour is critical to scalable analysis and control. Recent progress demonstrates its integration with modern methods. For example, Zheng et al. (2022) proposes a “top-down” control strategy for robotic swarms, utilising mean-field PDEs to model and control the swarm’s global density.

The approach generates velocity fields that guide the robots’ distribution to a desired target profile, with robustness guaranteed through Input-to-State Stability (ISS) analysis against perturbations and modelling errors. Although the velocity field design is centralised, the implementation on individual robots is fully distributed, making the solution practical for large-scale systems. This MFT-based control method is highly pertinent for resolving the common target problem in robotic swarms, as it optimises collective movement, positively impacting the arrival time at the destination and managing distribution to mitigate congestion and overcrowding.

In Dogbé (2010) the work is highly relevant to the application of MFT in robotic swarms. It proposes a framework for modelling the dynamics of pedestrian crowds using differential games, the dynamics of an individual pedestrian is described by the state vector $X(t) := (x, v)$, where x represents the position and v the velocity. The equations of motion are $\frac{d}{dt}x(t) = v(t)$ and $\frac{d}{dt}v(t) = \gamma(t)$, where $\gamma(t)$ is the acceleration vector that acts as the control.

The general cost functional for a trajectory, considering transport costs and a terminal cost, is given by

$$J(x, v) = \int_0^T L(t, x(t), v(t))dt + \phi(T, x(T)),$$

where $L(t, x, v)$ is the instantaneous cost function (running cost) and $\phi(T, x(T))$ is the terminal cost at the final position of the pedestrian at time T , reflecting the quality or effort to reach this final location. MFG is used to simplify the complex interactions among a large number of individuals, allowing each pedestrian's decisions to be modelled in relation to the statistical distribution of all others.

The Hamilton-Jacobi-Bellman (HJB) equation (2) below (Dogbé, 2010) is the mathematical apparatus governing the optimal dynamics of this agent continuum. Fundamentally, it describes how the optimal value (or minimum cost-to-go) for an individual agent changes over time and space, given its dynamics and the costs it incurs, serving as a cornerstone for optimal control problems in MFG; they emerge as a necessary condition for each agent to minimise its describes the optimal cost $u(t, x)$ for an individual pedestrian, given by

$$\frac{\partial u}{\partial t}(t, x) + H(x, \nabla u(t, x)) = 0, \text{ in } (0, T) \times \Omega, \quad (2)$$

with the boundary condition $u(T, x) = \phi(T, x)$. The Hamiltonian $H(x, p)$ is defined as $H(x, p) = \max_{v \in \mathbb{R}^d} \{-v \cdot p - L(x, v)\}$, where $p = \nabla u(t, x)$ is the momentum vector. The method derives a continuous formulation as PDEs for crowd dynamics, based on the minimisation of individual utility functions and the concept of Nash equilibrium (Karafyllis et al., 2010). This approach is pertinent to the common target problem in swarms, as it optimises collective movement and density distribution, directly impacting metrics.

The evolution of the pedestrian population density $\rho(t, x)$ is described by the Fokker-Planck equation (3) below (continuity equation), which describes how the probability $\rho(t, x)$ changes due to the flow of agents:

$$\frac{\partial \rho}{\partial t}(t, x) - \nabla \cdot (\rho(t, x) \bar{v}(t, x)) = 0, \text{ in } (0, T) \times \Omega, \quad (3)$$

where $\bar{v}(t, x)$ is the mean velocity field. In the context of MFG for pedestrian dynamics, the evolution of the density or distribution of a large population of agents over time and space, while accounting for their movement and ensuring conservation of the total number of agents, is modelled by a continuity equation. The core system of MFG Equations, which describes the Nash equilibrium when the number of pedestrians tends to infinity, consists of two coupled PDEs:

1. The HJB equation (4) below, for the optimal control problem of a representative agent, which now depends on the population density

$$\frac{\partial u}{\partial t}(t, x) + H(x, \nabla u(t, x), \rho(t, x)) = 0, \text{ in } (0, T) \times \Omega, \quad (4)$$

with the terminal condition $u(T, x) = \phi(x, \rho(T, x))$. The Hamiltonian $H(x, p, \rho)$ is defined as $H(x, p, \rho) = \max_{v \in \mathbb{R}^d} \{-v \cdot p - L(x, v, \rho)\}$. Here, the cost function $L(x, v, \rho)$ now depends on the population density ρ .

2. The Fokker-Planck equation (5) below represent the evolution of the population density, which depends on the value function u

$$\frac{\partial \rho}{\partial t}(t, x) - \nabla \cdot (\rho(t, x) \nabla_p H(x, \nabla u(t, x), \rho(t, x))) = 0, \text{ in } (0, T) \times \Omega, \quad (5)$$

with the initial condition $\rho(0, x) = \rho_0(x)$.

Having mapped the landscape, the continuum limit where PDEs replace discrete agents was zoomed into the last subsection.

Infinite-dimensional mean-field models

This section details the fundamental models extracted from the reviewed literature, which form the theoretical basis for the analysis of robotic swarm systems. The models, primarily derived from PDE, address congestion control and common target problems through mean-field approaches.

In these models, control generally focuses on managing the distribution of agents, which can have continuous or discrete states. It was assumed that the robots have local perception, allowing them to assess the density of robots in their neighbourhood or other aggregated properties of the swarm. Direct communication between robots is often reduced or absent, with interactions mediated by the mean field. These challenges have shifted attention towards the phenomena of diffusion, subdiffusion, and superdiffusion, which are recognised for introducing characteristics such as chaos and heightened dimensionality in robotic swarm simulations. This methodology effectively substitutes intricate, agent-specific dynamics with macroscopic approximations derived from population averages, commonly formulated through ordinary or partial differential equations such as the Fokker-Planck equation (5) (Auricchio et al., 2022).

Notably, Lévy flight within the realm of robotics, which emulates the foraging pathways of animals such as albatrosses, bumblebees, and deer, have been integrated into the depiction of complex dynamics. Beyond the basic paradigm of Brownian motion, characterized by Jean Perrin as the continuous, unpredictable motion of particles immersed in a liquid, Perrin articulated this phenomenon in French, stating: "Ils vont et viennent en tournoyant, montent, descendent, remontent encore, sans tendre aucunement vers le repos." This phrase describes the dynamic capacity of particles to move back and forth, as depicted two-dimensionally in Figure 2 and three-dimensionally in Figure 3.

Figure 2 - Two-dimensional projection of a single trajectory in two different times, highlighting collisions with fastmoving fluid molecules: (a) before, and (b) after.





Figure 3 - Three-dimensional representation of the stochastic displacement of microscopic particles suspended in a fluid.



Trajectory visualizations illustrate that Brownian motion follows a relatively compact and diffusive pathway, consisting of small, random steps, thereby maintaining proximity to its origin. Conversely, a Lévy flight displays intermittent, substantial leaps, allowing it to traverse greater distances from its point of origin within the same number of steps. This superdiffusive behavior is facilitated by the flight's heavy-tailed step-length distribution (a Lévy stable distribution), enabling more effective spatial exploration compared to Brownian motion, which is characterized by normally distributed step-lengths, resulting in a linear scaling of mean squared displacement over time. To model these complex phenomena, fractional calculus is utilised as a critical mathematical instrument. This facilitates the inclusion of memory effects and non-local interactions, thereby refining mean-field equations to more accurately depict anomalous diffusion and swarm exploration.

The fundamental equations presented by Estrada-Rodriguez and Gimperlein (2020) are crucial for modelling robotic swarm system dynamics, encompassing individual movement, stopping, and interactions. The microscopic movement of individuals is governed by the N -particle probability density function σ , which evolves according to the kinetic equations:

$$\frac{\partial \sigma}{\partial t} + \sum_{i=1}^N \left(\frac{\partial \sigma}{\partial \tau_i} + c \theta_i \cdot \nabla_{x_i} \sigma \right) = - \sum_{i=1}^N \beta_i \sigma,$$

where $\sigma = \sigma(x_1, \dots, x_N, t, \theta_1, \dots, \theta_N, \tau_1, \dots, \tau_N)$ is the N -particle probability density; $x_i \in \mathbb{R}^n$ is the position of robot i ; $\theta_i \in S = \{v \in \mathbb{R}^n : |v| = 1\}$ is its direction of motion; $\tau_i \geq 0$ is its elapsed run time since the last reorientation; $c > 0$ is the constant speed; β_i is the stopping frequency defined below. The probability that robot i has not stopped after run time τ_i follows the power law

$$\psi_i(x_i, \tau_i) = \left(\frac{\zeta_0(x_i)}{\zeta_0(x_i) + \tau_i} \right)^\alpha, \quad 1 < \alpha < 2,$$

with $\zeta_0(x_i) > 0$ a spatially-varying scale parameter. The stopping frequency is then

$$\beta_i(x_i, \tau_i) = -\frac{1}{\psi_i} \frac{\partial \psi_i}{\partial \tau_i} = \frac{\varphi_i(x_i, \tau_i)}{\psi_i(x_i, \tau_i)},$$

where $\varphi_i(x_i, \tau_i) = -\partial \psi_i / \partial \tau_i$, is the hazard rate. Upon stopping, a robot chooses a new direction θ_i^* via tumbling or alignment. The turning kernel is

$$T_i \phi(\theta_i^*) = \int_S k(x_i, t, \theta_i; \theta_i^*) \phi(\theta_i) d\theta_i,$$

where $k(x_i, t, \theta_i; \theta_i^*)$ is the transition probability density from direction θ_i to θ_i^* . The mean direction of the neighbours at position x_i is

$$\Lambda_i(x_i, \theta_i, t) = \frac{\mathcal{J}(x_i, t)}{|\mathcal{J}(x_i, t)|},$$

$$\mathcal{J}(x_i, t) = \int_{\mathbb{R}^n} \int_S K(|x_j - x_i|) p(x_j, t, \theta_j) \theta_j d\theta_j dx_j,$$

with $K(|x_j - x_i|)$ an influence kernel (decaying with distance); $p(x_j, t, \theta_j)$ the one-particle density of robots at (x_j, θ_j) . For the two-particle density $\sigma(x_1, x_2, t, \theta_1, \theta_2, \tau_1, \tau_2)$, the governing equation (6) is

$$\frac{\partial \sigma}{\partial \tau_1} + \frac{\partial \sigma}{\partial \tau_2} + \frac{\partial \sigma}{\partial t} + c \theta_1 \cdot \nabla_{x_1} \sigma + c \theta_2 \cdot \nabla_{x_2} \sigma = -(\beta_1 + \beta_2) \sigma. \quad (6)$$

After the robot 1 stops, its new run phase is initiated with

$$\tilde{\sigma}_{\tau_1}(x_1, x_2, t, \theta_1, \theta_2, \tau_1 = 0) = \int_S Q(\theta_1, \theta_1^*) \int_0^t \beta_1(x_1, \tau_1) \tilde{\sigma}_{\tau_1}(x_1, x_2, t, \theta_1^*, \theta_2, \tau_1) d\tau_1 d\theta_1^*,$$

where $\tilde{\sigma}_{\tau_1}$ denotes integration over τ_2 only; the reorientation operator is

$$Q(\theta_1, \theta_1^*) = \zeta k(x_1, t, \theta_1^*; \theta_1) + (1 - \zeta) \Phi(\Lambda_1 \cdot \theta_1),$$

with $\zeta \in [0, 1]$ the probability of tumbling (versus alignment); $\Phi(\cdot)$ the alignment distribution (depends on local mean direction Λ_1).

The work of H. Huang and Liu (2020) rigorously justifies the mean-field limit of an N -particle system, deriving the Vlasov-Poisson-Fokker-Planck (VPFP) (H. Huang & Liu, 2020) equations from a regularised microscopic N -particle system. The core of their work involves several key mathematical formulations that describe the particle dynamics at both microscopic and macroscopic levels along with specific interaction kernels. It follows that the microscopic N -particle system, subject to Brownian motions and interacting through a Newtonian potential in \mathbb{R}^3 , is described by the following system of Stochastic Differential Equations (SDE):

$$dx_i = v_i dt, \quad dv_i = \frac{1}{N-1} \sum_{j \neq i}^N k(x_i - x_j) dt + \sqrt{2\sigma} dB_i, \quad \text{for } i = 1, \dots, N. \quad (7)$$

The pairwise interaction between individuals, $k(x)$, is specifically modelled as the Coulombian kernel: $k(x) = a_{|x|^3}$. This microscopic system is approximated by the macroscopic VPFP equations for the probability density function $f(x, v, t)$ in phase space (x, v) :

$$\frac{\partial f(x, v, t)}{\partial t} + v \cdot \nabla_x f(x, v, t) + k * \rho(x, t) \cdot \nabla_v f(x, v, t) = \sigma \Delta_v f(x, v, t),$$

with $f(x, v, 0) = f_0(x, v)$. The charge density $\rho(x, t)$ within the VPFP system is defined as the integral of the probability density function over the velocity space $\rho(x, t) = \int_{\mathbb{R}^3} f(x, v, t) dv$.

The paper of H. Huang and Liu (2020) also considers a first-order stochastic system for particles:

$$dx_i = \frac{1}{N-1} \sum_{j \neq i}^N k(x_i - x_j) dt + \sqrt{2\sigma} dB_i,$$

for $i = 1, \dots, N$. The continuous description of the dynamics for this first-order system is given by the following non-linear PDE for the spatial density $f(x, t)$:

$$\frac{\partial f(x, t)}{\partial t} + \nabla \cdot (f(k * f)) = \sigma \Delta_x f.$$

Specific interaction kernels are presented for various applications. For fluid dynamics, the Biot-Savart kernel is given by:

$$k(x) = \frac{1}{2\pi} \left(\frac{-x_2}{|x|^2}, \frac{x_1}{|x|^2} \right).$$

For chemotaxis models, the Poisson kernel is often used. Thus, the paper introduces the notation for trajectories on phase space for the time evolution of the N -interacting Newtonian particles: $\Phi_t := (X_t, V_t) := (x_1^t, \dots, x_N^t, v_1^t, \dots, v_N^t)$, which explicitly states the Newtonian system with a noise term coupled to the velocity as a system of SDEs. It is essentially a detailed restatement of the equation (7) with explicit time dependence.

Furthermore, the equations of Auricchio et al. (2022) describe a novel Fokker-Planck-type model designed to simulate manufacturing processes through the collective dynamics of a large set of agents, aiming for uniform coverage of a target domain. The mesoscopic model in the study of Auricchio et al. (2022) translating the agent dynamics to a type of the same as in equation (5), with a constant diffusion and a discontinuous drift, expressed in divergence form as:

$$\frac{\partial f(x, t)}{\partial t} = \nabla_x \cdot [\psi(x, x_0) f(x, t) + \sigma^2 \nabla_x f(x, t)]. \quad (8)$$

In equation (8) below, the function characterising the drift is expressed by:

$$\psi(x, x_0) = \begin{cases} 0, & \text{if } |x - x_0| < \delta, \\ x - x_0, & \text{if } |x - x_0| \geq \delta, \end{cases} \quad (9)$$

where $D = \{x \in \mathbb{R}^d : |x - x_0| \leq \delta\}$ represents a d -dimensional sphere with centre x_0 and radius $\delta > 0$. The solution to the equation (8) is mass-preserving, meaning that for an initial probability density $f_0(x, t)$, the total mass remains conserved at any time $t \geq 0$: $\int_{\mathbb{R}^d} f(x, t) dx = 1$. The unique steady state of unit mass for the equation (8) is formally obtained by solving the differential equation (10) below

$$\sigma^2 \nabla_x f(x, t) = -\psi(x, x_0) f(x, t). \quad (10)$$

The solution to the equation (10) (the unique steady state $f^\infty(x)$) is given by:

$$f^\infty(x) = \begin{cases} \frac{m_1}{(2\pi\sigma^2)^{d/2}} \exp\left\{-\frac{(x-x_0)^2}{2\sigma^2}\right\}, & \text{if } |x - x_0| \geq \delta, \\ m_2 \left(\frac{\delta^d \pi^{d/2}}{\Gamma(d/2+1)}\right)^{-1}, & \text{if } |x - x_0| < \delta. \end{cases} \quad (11)$$

The same equilibrium configuration as in the equation (11) can also be obtained from an alternative formula similar to the equation (12) below, with a variable diffusion coefficient:

$$\frac{\partial f(x, t)}{\partial t} = \nabla_x \cdot [(x - x_0)f(x, t) + \nabla_x(\kappa(x)f(x, t))]. \quad (12)$$

In the equation (12), the continuous function characterising the diffusion coefficient $\kappa(x)$ is expressed by

$$\kappa(x) = \begin{cases} \sigma^2 + \frac{\delta^2}{2} - \frac{1}{2}|x - x_0|^2, & \text{if } |x - x_0| < \delta, \\ \sigma^2, & \text{if } |x - x_0| \geq \delta. \end{cases} \quad (13)$$

The values m_1 and m_2 , which quantify the percentages of mass for the Gaussian and uniform parts, respectively, are determined by solving the following linear system resulting from mass conservation and continuity conditions:

$$\begin{cases} m_1(1 - \text{erf}(\frac{\delta}{\sqrt{2\sigma^2}})) + m_2 = 1, \\ \frac{m_1}{\sqrt{2\pi\sigma^2}} \exp\left\{-\frac{\delta^2}{2\sigma^2}\right\} - \frac{m_2}{2\delta} = 0. \end{cases} \quad (14)$$

Similarly, M. Huang et al. (2006) addresses stochastic dynamic games in large populations where agents are weakly coupled, using the Nash Certainty Equivalence (NCE) Principle, to derive decentralised control synthesis via McKean-Vlasov systems. The individual agent optimal control problem is defined by minimising a cost functional J , considering the population distribution P_s :

$$J(t, x, v, P) = \inf_{u_t} E_{t,x,v} \left[\int_t^\infty r(s, X_s, v_s, P_s) ds + S(X_s) \right].$$

The optimal value function $V(t, x, v, P)$ for the representative agent is governed by the equation (15) below:

$$\frac{\partial V}{\partial t} + \max_u [L(t, x, v, u, P) + \nabla_x V(t, x, v, P) \cdot f(t, x, v, u, P)] + \text{Tr} [D \nabla_v^2 V(t, x, v, P)] = 0. \quad (15)$$

The evolution of the population distribution $P(t, x, v)$ is described by the equation (16) below, which is of Fokker-Planck type:

$$dP(t, x, v) = [-\nabla_x \cdot (f(t, x, v, \bar{u}(t, x, v, P))P(t, x, v))] dt + \text{Tr}[D \nabla_v^2 P(t, x, v)] dt. \quad (16)$$

The usage of the overbar observed with the Fokker-Planck type equation means an averaged quantity, as in the term $\bar{u}(t, x, v, P)$. Here, \bar{u} denotes a collective or spatially averaged quantity, often referred to as the mean field. This structure is critical because it introduces non-linearity into the system dynamics: the local transport (advection) of particles is not only governed by local states (x, v) but is also influenced by the aggregated, mean behaviour of the entire population distribution P . This dependency is fundamental to models describing self-organising systems and collective dynamics.

In the specific context of a Linear Quadratic Gaussian (LQG) framework, the individual agent dynamics is given by the SDE:

$$dx_t = (A_0 x_t + A_1 \bar{x}_t + B u_t) dt + F dw_t.$$

The quadratic cost functional for each individual agent in the LQG problem is

$$J(u) = E \left[\int_0^\infty (\|x_t\|_Q^2 + \|u_t\|_R^2) dt \right].$$

The HJB equation (15) for the value function $V(x)$ in this LQG setting becomes the quadratic equation (17) below:

$$\frac{1}{2}x^T Qx + \frac{1}{2}\bar{x}^T \bar{Q}\bar{x} + \nabla V(x)^T (A_0 x + A_1 \bar{x}) + \text{Tr}[F^T F \nabla^2 V(x)] - \frac{1}{2} \nabla V(x)^T B R^{-1} B^T \nabla V(x) = 0. \quad (17)$$

In the stationary HJB equation for the Linear-Quadratic-Gaussian (LQG) control problem, the terms \bar{x} and \bar{Q} denote quantities linked to state uncertainty. \bar{x} signifies the conditional mean state, the optimal estimate of the true system state x , typically supplied by a Kalman filter. Its presence in the quadratic cost function ($\bar{x}^T \bar{Q} \bar{x}$) and the drift term demonstrates the fundamental separation principle of LQG control. Furthermore, \bar{Q} is the cost weighting matrix specifically associated with the estimated state \bar{x} . While Q penalises the deviation of the true state x , \bar{Q} allows the cost function to also penalise or reward outcomes based on the estimated or perceived state, which is crucial for systems operating under partial observation.

From the equation (17), the optimal control law u_t for the individual agent is derived as $u_t = -R^{-1} B^T \nabla V(x_t)$. Applying this optimal control yields the linearised state equation (18) below for x_t :

$$dx_t = (A_0 x_t + A_1 \bar{x}_t - B R^{-1} B^T P x_t) dt + F dw_t. \quad (18)$$

In the equation (18), \bar{x}_t denotes the estimated state at time t . This equation describes the evolution of the actual state x_t when subject to the optimal control policy. Although the system state is x_t , the optimal control input is calculated based only on the best available estimate, \bar{x}_t , because the true state is corrupted by noise. The appearance of \bar{x}_t in the drift term, $A_1 \bar{x}_t$, therefore explicitly confirms that the controller is basing its corrective action on its filtered knowledge of the system, a key feature of optimal stochastic feedback control.

The resulting McKean-Vlasov SDE for the individual state x_t , incorporating the mean-field effect ($E[x_t]$), is

$$dx_t = (A_0 x_t + A_1 E[x_t] - B R^{-1} B^T P x_t) dt + F dw_t.$$

Thus, the evolution of the mean of the state $E[x_t]$ is governed by a linear Ordinary Differential Equation (ODE) as in the equation (19) below:

$$\frac{d}{dt} E[x_t] = (A_0 + A_1 - B R^{-1} B^T P) E[x_t]. \quad (19)$$

Another study related to infinite-dimensional mean-field model is the MFG concept. Its primary objective, as introduced by Lasry and Lions (2007), is to tackle the formidable complexity found in systems that involve an exceptionally large number of rational agents. This approach seeks to make such intricate systems manageable for analysis and modelling.

MFG theory eschews the computationally intractable task of modelling every pairwise interaction, a problem of exponential complexity, in favour of a paradigm shift: it models the decision-making of a single, representative agent. This agent does not react to specific counterparts but instead responds to the aggregate and anonymised influence of the entire population, termed the Mean-Field.

This innovative strategy reduces an ostensibly high-dimensional problem to a coupled system of two PDEs: one governing the individual optimal strategy, equation (20), and the other describing the temporal evolution of the population distribution, equation (21). Thus, Lasry and Lions (2007) seek to establish and validate a rigorous yet tractable mathematical framework capable of identifying equilibrium solutions in complex collective dynamics that would otherwise be analytically unattainable from the equations

$$-\frac{\partial u}{\partial t} + H(x, Du) + F(x, m) = 0 \text{ in } (0, T) \times \mathbb{T}^d, u(T, x) = G(x) \text{ in } \mathbb{T}^d, \quad (20)$$

$$\frac{\partial m}{\partial t} - \nabla_x \cdot (m H_p(x, Du)) = 0 \text{ in } (0, T) \times \mathbb{T}^d, m(0, x) = m_0(x) \text{ in } \mathbb{T}^d. \quad (21)$$

Here, $u(t, x)$ is the value function (or function of the cost to go) for a representative agent, and $m(t, x)$ is the density of agents in the state space. The time horizon is denoted by T and \mathbb{T}^d (a d -dimensional torus) as the spatial domain is considered periodic.

The Hamiltonian function $H(x, p)$ describes the dynamics of the individual agent and the cost associated with its control:

$$H(x, p) = \frac{1}{2}|p|^2. \quad (22)$$

The optimal control $\alpha^*(t, x)$ for the representative agent is determined by the first-order condition $\alpha^*(t, x) = -Du(t, x)$. By substituting the specific Hamiltonian equation (22) into the general MFG system equation (20) and equation (21), the HJB equation (23) below for the optimal control of a single agent takes the form

$$-\frac{\partial u(t, x)}{\partial t} + \frac{1}{2}|Du(t, x)|^2 + F(x, m(t, x)) = 0, \quad (23)$$

and the density of agents, which is consistent with the optimal individual dynamics, is determined by the equation (24) below:

$$\frac{\partial m(t, x)}{\partial t} - \nabla_x \cdot (m(t, x)(-Du(t, x))) = 0. \quad (24)$$

Lasry and Lions (2007) also discuss two advanced mathematical concepts in control theory: classical optimal control and MFG. Classical optimal control uses the coupling of state PDEs with backward-propagating adjoint equations to optimise the system trajectory. In strategic settings, MFG refines this approach by seeking equilibrium through coupled PDE systems. The text emphasises the importance of robustness alongside optimality, highlighting the role of Input-to-State Stability (ISS), which ensures system stability and bounded error amid external disturbances like sensor noise or model inaccuracies. These methods, while distinct in their goals – optimality and robustness – utilise mean-field abstraction to derive control laws that ensure efficiency, scalability, and reliability in complex systems.

Zheng et al. (2022) describe the evolution of the swarm's mean-field density $\rho(x, t)$ over a bounded spatial domain $\Omega \subset \mathbb{R}^d$ by the mean-field PDE

$$\frac{\partial \rho}{\partial t}(x, t) = \nabla \cdot (A(x)\rho(x, t)) + \frac{\sigma^2}{2}\Delta\rho(x, t) + B(x)u(x, t)\rho(x, t), \quad (25)$$

where $A(x)$ and $B(x)$ are matrices describing the drift and control input distribution, respectively, σ^2 is the diffusion coefficient, and $u(x, t)$ is the velocity field control input. The control objective is to steer the swarm density from an initial profile $\rho_0(x)$ to a desired target profile $\rho_f(x)$ at time T . This is achieved by minimising the following objective function

$$J(u) = \int_0^T \int_{\Omega} \|u(x, t)\|_2^2 dx dt + \int_{\Omega} \|\rho(x, T) - \rho_f(x)\|_2^2 dx. \quad (26)$$

The optimal control problem for this system is solved using the equation (23). The HJB equation (27) below represents the value function $V(x, t)$ (which is related to the optimal cost to go) is given by:

$$\frac{\partial V}{\partial t}(x, t) = \nabla V(x, t) \cdot A(x) + \frac{\sigma^2}{2}\Delta V(x, t) - \frac{1}{2}B(x)^T B(x)(\nabla V(x, t))^2 + \lambda(x)\rho(x, t). \quad (27)$$

The optimal control law $u(x, t)$ that minimises the objective function equation (26) is derived from the equation (23) as a feedback law

$$u(x, t) = -B(x)^T \nabla V(x, t). \quad (28)$$

The full characterisation of the optimal solution also requires the adjoint (costate) equation for $\lambda(x, t)$, that is, the Lagrange multipliers corresponding to the constraints of the minimisation of the cost in the equation (26):

$$\frac{\partial \lambda}{\partial t}(x, t) = -\nabla \cdot (A(x)\lambda(x, t)) - \frac{\sigma^2}{2}\Delta\lambda(x, t) + \lambda(x, t)B(x)^T B(x)\nabla V(x, t). \quad (29)$$

By substituting the optimal control equation (28) back into the swarm dynamics equation (25), the total velocity field $v(x, t)$ of the swarm can be expressed as

$$v(x, t) = -A(x) - \sigma^2 \frac{\nabla \rho(x, t)}{\rho(x, t)} - B(x)B(x)^T \nabla V(x, t).$$

The paper of Zheng et al. (2022) also leverages the concept of Input-to-State Stability (ISS) to show that the perturbed closed-loop system is stable.

Moreover, Ornia et al. (2022) explores how MFT can be used to reformulate a stochastic multi-agent foraging problem into a deterministic autonomous system, enabling the analysis of limit behaviours and optimality guarantees. The individual agent state evolution (position X_t) in the multi-agent system is described by the SDE

$$dX_t = (v_c + \beta_d(C(X_t)) - \gamma_c C(X_t))dt + \sigma dW_t. \quad (30)$$

Here, v_c is a constant velocity, $\beta_d(C(X_t))$ represents a density-dependent velocity, $\gamma_c C(X_t)$ is a consumption or decay term, and W_t is a Wiener process with a diffusion coefficient σ . The pheromone concentration $C(x)$ at a location $x \in \mathbb{T}^d$ is a function of the positions of the agents. In the discrete N -agent system, it is defined as

$$C(x) = \frac{1}{N} \sum_{i=1}^N \frac{1}{\epsilon^d} \chi_{B_\epsilon(x_i)}(x),$$

where $\chi_{B_\epsilon(x_i)}(x)$ is the characteristic function for a ball of radius ϵ centred at x_i . The macroscopic behaviour of the agent swarm is described by the equation (31) below for the agent density $m(t, x)$:

$$\frac{\partial m(t, x)}{\partial t} + \nabla \cdot (m(t, x)V(C(x))) - \frac{\sigma^2}{2} \Delta m(t, x) = 0, \quad (31)$$

where $V(C(x))$ is the velocity field determined by the pheromone concentration. The evolution of the pheromone field itself is modelled by $\frac{\partial C(x)}{\partial t} = \alpha m(t, x) - \lambda C(x)$.

Here, α is the deposition rate of pheromone by agents, and λ is the decay rate of pheromone.

The deterministic part of the velocity field $V(C(x))$ for the agents which is influenced by the pheromone concentration $C(x)$, derived from the SDE in the equation (30), is expressed by

$$V(C(x)) = v_c + \beta_d C(x) - \gamma_c C(x). \quad (32)$$

In essence, the equation (32) models the agents' velocity field as a combination of a constant base velocity, a velocity component that responds positively to pheromone concentration, and another component that acts as a decay or consumption factor related to the pheromone concentration. The paper then proceeds to analyse the steady-state behaviour of the system. In the steady state, the time derivatives are zero, leading to a coupled system for the stationary density $m(x)$ and concentration $C(x)$:

$$\nabla \cdot (m(x)V(C(x))) - \frac{\sigma^2}{2} \Delta m(x) = 0, \quad (33)$$

$$\alpha m(x) - \lambda C(x) = 0. \quad (34)$$

From the equation (34), a direct relationship between the steady-state density and concentration is found: $m(x) = \frac{\lambda}{\alpha} C(x)$. Substituting this relation into equation (33) yields a single differential equation for the steady-state concentration $C(x)$:

$$\nabla \cdot (C(x)V(C(x))) - \frac{\sigma^2}{2} \Delta C(x) = 0. \quad (35)$$

For the one-dimensional case, the equation (35) simplifies to an ODE:

$$\frac{d}{dx} (C(x)V(C(x))) - \frac{\sigma^2}{2} \frac{d^2}{dx^2} C(x) = 0.$$

The implementation of MFT in robotic swarms relies not on a singular model, but rather on a diverse range of mathematical tools selected according to the task's intrinsic characteristics. Such adaptability is essential, as it manages the otherwise insurmountable complexity of multi-agent systems by substituting pairwise interactions with the effect of an anonymous mean field exerted on a representative agent. In continuous spatial settings, this abstraction manifests as a PDE, and the careful selection of the governing PDE is the initial decisive step toward effective control. In scenarios where agent motion is predominantly deterministic, such as formation flight, the continuity equation, an example is equation (3), is identified as the appropriate mechanism, determining the evolution of the swarm's numerical density $\rho(x, t)$.

By contrast, in environments characterised by stochasticity or inherent uncertainty, such as exploratory missions, the equation (31) becomes essential, characterising the evolution of the probability density $w(x, t)$ (Carmona & Delarue, 2018) of an agent presence in a given state.

This distinction is not merely descriptive; it underpins the development of advanced control strategies. The resulting PDE, in whatever form it may take, serves as the basis for determining an optimal trajectory through adjoint-based optimisation or ensuring robustness against real-world disturbances via Input-to-State Stability (ISS) analysis.

Hence, the ability to select the appropriate mean-field model enables MFT to not only represent, but also regulate a broad spectrum of collective behaviours in a scalable and reliable manner.

Finite-dimensional mean-field approximations

In Elamvazhuthi and Berman (2019), the authors propose a dual-faceted modelling approach wherein finite-dimensional systems are employed to capture low-order statistics, and infinite-dimensional limits are utilised to ensure precision for extensive populations. Both continuous-time and discrete-time frameworks are considered simultaneously, resulting in a well-structured hierarchy of representations.

Continuous-time

The macroscopic evolution of an ensemble of agents, also known as Kolmogorov forward equation, is encoded in the probability-mass vector $x(t) \in \mathcal{P}(\mathcal{V})$, whose components obey:

$$\dot{x}(t) = \sum_{e \in \mathcal{E}} u_e(t) B_e, \quad x(0) = x^0. \quad (36)$$

Here, $\dot{x}(t)$ denotes the temporal rate of change of the probabilities; each entry $x(t)$ is the distribution of the random variables and $\mathbb{P}(X_i(t) = v) = x_v(t)$ is the likelihood that an arbitrary agent occupies state v at time t that determines the time evolution of the probability densities.

The scalar control parameter $u_e(t)$ represents the transition rate along the directed edge e , while the matrix $B_e \in \mathbb{R}^{M \times M}$ encodes the net effect of that transition on the probability vector: its entries satisfy $B_e^{ij} = -1$ when $i = j$ is the source vertex of e , $B_e^{ij} = +1$ when i is the target vertex and j the source, and vanish otherwise. Consider an example with a three-state Markov chain. In this case, this construction is given by

$$\dot{x}_1(t) = -u_{12}(t)x_1(t) + u_{21}(t)x_2(t), \quad (37)$$

$$\dot{x}_2(t) = -(u_{21}(t) + u_{23}(t))x_2(t) + u_{12}(t)x_1(t) + u_{32}(t)x_3(t), \quad (38)$$

$$\dot{x}_3(t) = -u_{32}(t)x_3(t) + u_{23}(t)x_2(t), \quad (39)$$

with prescribed initial masses $x_i(0) = x_i^0$. To steer the population towards a prescribed distribution, transition rates are closed in feedback form with a mean-field feedback, whose laws can describe the mechanism of quorum sensing in biological swarms such as bacterial colonies, whereby individuals' assessment of population density triggers a change in their behaviour and gives rise to coordinated collective phenomena. This feedback is given by $k_e(x) = k_e^* + \sigma_{S(e)}(x_{S(e)}, q_{S(e)})(\alpha - 1)k_e^*$, where k_e^* is a baseline rate, $q_{S(e)}$ denotes a target quorum fraction at vertex $S(e)$. The analytical switching function

$$\sigma_{S(e)}(x) = \left(1 + \exp \left(\gamma \left(q_{S(e)} - \frac{x_{S(e)}}{x_{S(e)}^d} \right) \right) \right)^{-1} \in [0, 1],$$

modulates the transition rate according to the current local density, where $\gamma > 0$ denotes the responsiveness parameter controlling the steepness of the switching behaviour, $q_{S(e)}$ represents the threshold parameter for switching activation, and $x_{S(e)}^d$ signifies the desired target density for state $S(e)$. The parameter γ tunes the system's responsiveness to density deviations, whilst the threshold $q_{S(e)}$ determines the saturation characteristics of the switching mechanism.

Elamvazhuthi and Berman (2019) rigorously justify the deterministic approximation by establishing the mean-field limit when approximating large finite-dimensional systems with simpler infinite-dimensional models. Specifically, when individual transition rates adopt the form

$$u_e(t) = v_e \left(\frac{1}{N} \sum_{i=1}^N \chi_1(X_i(t)), \dots, \frac{1}{N} \sum_{i=1}^N \chi_M(X_i(t)) \right), \quad (40)$$

where N denotes the total number of agents, $X_i(t)$ represents the state of the i -th individual agent at time t , χ_j are characteristic functions for $j = 1, \dots, M$, and M is the number of discrete states, the empirical measure $\bar{Y}_N(t) = \left[\frac{1}{N} \sum_{i=1}^N \chi_1(X_i(t)), \dots, \frac{1}{N} \sum_{i=1}^N \chi_M(X_i(t)) \right]^T$ converges almost surely to the solution of the limiting system $\dot{x}(t) = \sum_{e \in E} v_e(x_1, \dots, x_M) B_e x(t)$, $x(0) = x^0$, where v_e are transition rate functions, B_e represent transition operators, and x^0 denotes the initial condition vector. The convergence holds in the sense that $\lim_{N \rightarrow \infty} \sup_{t \geq 0} \|\bar{Y}_N(t) - x(t)\| = 0$. This result legitimises the low-dimensional ODE as an exact descriptor for large ensembles.

Generalising the feedback control law, Bensoussan et al. (2018) proposes an arbitrary distribution stabilisation approach through $u_e(t) = f_e(x_{S(e)}(t))$, where f_e is any continuously differentiable, non-decreasing mapping satisfying $f_e(y) = 0$ if and only if $y = x_{S(e)}^d$. This construction guarantees global asymptotic convergence to the desired macroscopic profile.

To encompass mass-action kinetics, the mean-field representation extends naturally to chemical reaction networks through

$$\dot{x}(t) = \sum_{r \in R} k_r(t) f_r(x(t)), \quad x(0) = x^0 \in \mathbb{R}_{\geq 0}^M, \quad (41)$$

where R denotes the set of reactions, $k_r(t)$ represents the time-dependent rate constant of reaction r , and f_r signifies the associated stoichiometric vector field.

The mathematical structure of the equation (41) mirrors that of the population dynamics system, both sharing the fundamental form $\dot{x}(t) = \sum [\text{rate functions} \times \text{state-dependent terms}]$. This structural similarity underscores the unifying power of the mean-field perspective, demonstrating how this theoretical framework provides a common mathematical foundation for describing collective behaviour across diverse physical, biological and chemical systems.

Discrete-time

When agent states evolve according to a Markov chain, the mean-field evolution is captured by the difference equation (42) below, proposed by Bensoussan et al. (2018):

$$x(n+1) = \sum_{e \in \mathcal{E}} u_e(n) B_e x(n), \quad x(0) = x^0 \in \mathcal{P}(\mathcal{V}), \quad (42)$$

where $x(n) \in \mathbb{R}^M$ is the probability vector at step n , the scalar $u_e(n) \in [0, 1]$ denotes the transition probability along the edge e , and the matrix B_e satisfies $B_e^{ij} = 1$ when i is the target vertex and j the source of e , and vanishes otherwise. This discrete analogue mirrors its continuous counterpart and enables unified analysis across sampling regimes.

In summary, discrete-time systems exhibit significant advantages over their continuous-time counterparts within the realm of contemporary technological applications. These advantages are primarily attributed to their exceptional computational efficiency, inherent numerical robustness, and fundamental alignment with modern digital architectures.

While continuous systems require complex and potentially unstable numerical integration, discrete models function based on straightforward algebraic principles, thereby offering enhanced stability and predictability.

Although constraints such as aliasing necessitate meticulous design consideration, they are both well-understood and quantifiable. Consequently, the adoption of discrete-time frameworks signifies a pivotal and irreversible paradigm shift, establishing them as the essential cornerstone for future advancements in engineering and the applied sciences.

Mean-field models from statistical physics

The seminal work by Thouless et al. (1977) rectified inconsistencies in the Sherrington-Kirkpatrick “spin glass” model at low temperatures by developing an MFT that considers both the average agent state and its fluctuations, resulting in a thermodynamically stable description. For robotic swarms, this implies the need for mean-field models that capture the distribution and heterogeneity of agent states, not just their average behaviour.

Bonnet (2019) presents a rigorous mathematical framework for the optimal control of multi-agent systems in the mean-field limit, employing optimal transport theory. The central innovation lies in generalising Pontryagin’s Maximum Principle to the Wasserstein space, providing the necessary conditions to determine controls that optimise the behaviour of an agent density.

The work by Bonnet (2019) establishes the conditions for the spatial regularity of optimal controls, a crucial step to ensure the applicability of MFC to systems with a finite number of agents (microscopic). It also explores control strategies to guarantee alignment (flocking) in cooperative systems, even in the face of communication failures, which has direct application in developing robust algorithms for robotic swarms. That work develops a complex mathematical structure, with its most representative formulae defining the models at different scales (micro and macro) and the optimal control problem.

The microscopic model (N -agents) is described by the controlled N -agent dynamics, which expresses how the dynamics of each agent i are influenced by the other agents and by an individual control $u_i(t)$:

$$\frac{d}{dt}(x_i(t)) = v^N[x(t)](t, x_i(t)) + u_i(t), \quad (43)$$

where x_i is the state of an agent i , $x(t)$ is the state vector of all N agents, v^N is the interaction velocity field, and $u_i(t)$ is the applied control.

The Macroscopic Model (Mean-Field Limit) is given by the controlled continuity equation (44) below, which describes the evolution of the probability density $\mu(t)$ of finding an agent in a given state, in the limit as $N \rightarrow \infty$:

$$\frac{\partial \mu(t)}{\partial t} + \nabla \cdot ((v[\mu(t)](t, \cdot) + u(t, \cdot))\mu(t)) = 0, \quad (44)$$

where $\mu(t)$ is the probability measure of the agent distribution, $v[\mu(t)]$ is the non-local interaction velocity field (dependent on μ), and $u(t, \cdot)$ is the MFC.

The Optimal Control Problem in Wasserstein Space is formulated as the General Optimisation Problem (GOP), seeking to minimise a cost subject to the continuity equation (45) below:

$$\min_{u \in U} \left[\int_0^T L(t, \mu(t), u(t)) dt + \phi(\mu(T)) \right] \quad (45)$$

subject to $\frac{\partial \mu(t)}{\partial t} + \nabla \cdot ((v[\mu(t)] + u(t, \cdot))\mu(t)) = 0$ with $g_i(\mu(t)) \leq 0$ and $h_j(\mu(t)) = 0$,

where L is the running cost, ϕ is the final cost, and g_i, h_j are the state constraints on μ . The density μ is a measure over the spatial domain and $\mu(t)$ denotes the entire spatial distribution at time t . It implicitly contains information about the spatial variable. For example, $\mu(t)$ could be written as $d\mu(t)(x)$, where x is the spatial variable. In contrast, $u(t, \cdot)$ explicitly shows its dependence on both time t and space x . The control u is typically a velocity field that varies in time and across the spatial domain to influence the movement of agents at different locations.

The optimality condition (Pontryagin’s maximum principle) is expressed by the Forward-Backward Hamiltonian System. This central result of Bonnet (2019) generalises Pontryagin’s maximum principle, stating that an optimal trajectory and control ($\mu^*(t)$ and $u^*(t)$) must satisfy a system of equations for the state and co-state, along with a maximisation condition.

Bonnet (2019) Research in robotic swarms has been extended to encompass advanced algorithmic methodologies which are aimed at optimising planning and decision making within multi-agent systems. The experimental validation of mean-field theory in robotic swarms has been demonstrated across a variety of collective tasks.

Through the careful comparison of quantitative metrics, such as cost and efficiency rates, spatial distributions, and task performance, researchers have shown that mean-field models can accurately predict the emergent behaviour of real-world robotic swarms. For a comprehensive comparison of the models, the Table 2 has been created within this systematic literature review, structured according to the following key criteria.

Table 2 - Systematic Comparison of Mean-Field Swarm Robotic Models

| Feature | HJB/FP Models (Mathematical Components) | Physics-Informed Emergent Models (e.g., Spin/Vicsek/Kuramoto) | MFC/MFG Frameworks (MFT Control/Equilibrium) |
|---------------------------------|---|---|--|
| Primary Objective | HJB: Find the optimal control policy u^* for individual agents (stochastic optimisation). FP: Model and control the evolution of the probability density $\rho(x, t)$. | Describe and explain <i>emergent</i> collective behaviours based on local rules; analyse phase transitions and self-organisation. | MFC: Provide an optimal control law for the <i>density distribution</i> ρ (coupled global optimisation). MFG: Find the <i>Nash Equilibrium</i> (optimal individual strategies \leftrightarrow consistent population density). |
| Modelling Scale | HJB: Microscopic (focus on individual path $x(t)$). FP: Macroscopic (focus on density $\rho(x, t)$). | Microscopically defined (local interactions) with macroscopically emergent outcomes (e.g., global polarisation). | Intrinsically macroscopic (treats the swarm as a continuous density in the limit $N \rightarrow \infty$). |
| Nature of Solution | HJB: A “value function” $V(x, t)$ (optimal cost). FP: A probability distribution function $p(x, t)$. | An order parameter characterising the system’s phase state (e.g., polarisation, average magnetisation). | MFC: Optimal control law u^* for the collective system. MFG: The coupled, self-consistent pair (V^*, ρ^*) (solutions to the coupled HJB and FP PDEs). |
| Computational Complexity | HJB: High; suffers from the “curse of dimensionality” for finite N . FP: Requires solving PDEs (scales better with N). | Direct (agent-based) simulations are scalable $O(N)$; equilibrium analysis can be NP-hard. | Requires solving coupled, non-linear PDEs (HJB and FP) in the continuous space; high analytical complexity. |
| Information Requirements | HJB: Precise dynamical model, cost functional. FP: Interaction rules and diffusion/noise function. | Only requires the definition of local interaction rules and a noise/temperature parameter. | Requires individual agent cost functions and interaction kernels, assuming agents respond to the mean field ρ . |
| Control Type | HJB: Centralised (in finite systems). FP: Density control, decentralisable. | Inherently decentralised (based on local neighbourhood interactions). | MFC: Global optimisation (centralised). MFG: Decentralised strategy (agents react to the mean field ρ , not specific individuals). |
| Typical Application | HJB: Optimal trajectory planning for small groups. FP: Area coverage, density management, congestion control. | Collective decision-making (quorum sensing), aggregation, pattern formation, stability analysis. | Large-scale congestion control, resource optimisation (MFC), strategic and competitive interactions (MFG). |

The Table 3 serves to systematically delineate the core conceptual and mathematical differences between the two primary MFT frameworks: Mean Field Control (MFC) and Mean Field Game (MFG), it maps the distinct primary focus and solution processes of each, enabling a clear taxonomic classification of the extant literature on robotic swarms and providing the necessary theoretical foundation for discussing how each framework contributes to the estimation of performance metrics in target acquisition scenarios.

Table 3 - Structure of the Coupled Mean Field Control and Mean Field Game Systems

| Framework | Primary Focus | Optimisation Eq. (HJB-like) | Dynamics Eq. (FP-like) | Solution Process |
|----------------------|------------------------------------|------------------------------------|------------------------|--|
| MFC (Control) | Global cost minimisation $J(\rho)$ | Variational (optimality condition) | Controlled FP | Find the control law u^* that minimises $J(\rho^*)$, where ρ^* is the resulting distribution under that law. |
| MFG (Game) | Nash equilibrium | HJB (Hamilton–Jacobi–Bellman) | FP (Fokker–Planck) | Fixed-point solution: find the pair (V^*, ρ^*) that satisfies both individual optimality and population consistency simultaneously. |

Conclusions

Drawing together the ideas above, the applicability of MFT as a foundational methodology for the scalable estimation of performance in robotic swarms was investigated. This review of the literature demonstrated that MFT provides a robust and promising mathematical framework for addressing the challenges of high-dimensionality and complex interactions in multi-agent systems.

A substantial increase in scholarly investigation was noted in this area, particularly from 2010 onwards, indicating MFT's growing significance and potential. Among the various applications explored, those related to optimal control, coordination, and emergent behaviour in robotic systems were particularly prominent. A recurring conclusion is MFT's capability to simplify the analysis of complex systems, facilitating the derivation of collective behaviours from individual interaction rules.

However, interpreting MFT outcomes in robotic swarms needs a meticulous assessment of the extent to which the predicted collective behaviour corresponds with empirical observations. Discrepancies may arise from the disparity between idealised models and the intricacy of real-world systems.

While MFT's principal advantage resides in its scalability (where the computational effort per robot ideally remains unaffected by swarm size), several inherent limitations are encountered by the method. These include: the assumption of a large number of agents N , which results in a diminished precision for small swarms due to significant stochastic fluctuations; the presumption of agent homogeneity in classical models, contrasting with observed variations in reality; the nature of the mean approximation, which tends to smooth out critical fluctuations and the effect of outliers; the mathematical complexity intrinsic to the MFT model, often manifesting as a non-linear PDE-based or an MFG system, posing difficulties for real-time resolution; and the incongruity between theoretical global mean-field interactions and practical local interactions, necessitating theoretical justification for how locality impacts model validity.

Despite these challenges, the strengths identified, such as the ability to simulate large-scale systems with computational efficiency, are indisputable. The insights acquired will subsequently be used from this review to estimate key performance metrics including total simulation time, target arrival time, and target departure time, for a swarm algorithm based on potential fields as those developed by Passos et al. (2023). These estimates will be validated against outcomes from direct simulations, aiming to enhance the existing state-of-the-art and advance the practical application of MFT in robotic systems.

Acknowledgments

The authors thank the Federal University of Recôncavo da Bahia (UFRB), FAPESB and CNPq for their support and studentships.

Author contributions

J. V. N. de S. Nunes contributed to the Conceptualisation, Investigation, Formal Analysis, Visualisation, Writing and Original Draft. **Y. T. dos Passos** contributed to the Conceptualisation, Methodology, Supervision, Writing and Revision.

Conflicts of interest

The authors declare that no competing financial or commercial interests exist in relation to the work described in this manuscript.

References

- Auricchio, F., Toscani, G., & Zanella, M. (2022). Fokker-Planck Modeling of Many-Agent Systems in Swarm Manufacturing: Asymptotic Analysis and Numerical Results. *ResearchGate*. <https://doi.org/10.13140/RG.2.2.25950.72006>
- Bensoussan, A., Huang, T., & Laurière, M. (2018). Mean field control and mean field game models with several populations. *arXiv*, 2, 1–32. <https://doi.org/10.48550/arXiv.1810.00783>
- Bensoussan, A., Frehse, J., & Yam, P. (2013). *Mean Field Games and Mean Field Type Control Theory*. Springer. <https://doi.org/10.1007/978-1-4614-8508-7>
- Bonabeau, E., Dorigo, M., & Theraulaz, G. (1999). *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press. <https://doi.org/10.1086/393972>
- Bonnet, B. (2019). *Optimal Control in Wasserstein Spaces*. [Theses, Aix-Marseille Université]. <https://hal.science/tel-02361353>
- Burger, M., Di Francesco, M., Markowich, P. A., & Wolfram, M.-T. (2014). Mean Field Games with Nonlinear Mobilities in Pedestrian Dynamics. *Discrete and Continuous Dynamical Systems - B*, 19(5), 1311–1333. <https://doi.org/10.3934/dcdsb.2014.19.1311>
- Carmona, R., & Delarue, F. (2018). *Probabilistic Theory of Mean Field Games with Applications I-II*. Springer. <https://doi.org/10.1007/978-3-319-58920-6>
- Couzin, I. D., Krause, J., James, R., Ruxton, G. D., & Franks, N. R. (2002). Collective Memory and Spatial Sorting in Animal Groups. *Journal of Theoretical Biology*, 218(1), 1–35. <https://doi.org/10.1006/jtbi.2002.3065>
- Cui, K., Li, M., Fabian, C., & Koepl, H. (2023). Scalable task-driven robotic swarm control via collision avoidance and learning mean-field control. In Institute of Electrical and Electronics Engineers, 2023 IEEE International Conference on Robotics and Automation (ICRA) [Proceedings]. International Conference on Robotics and Automation, London, United Kingdom. <https://doi.org/10.1109/ICRA48891.2023.10161498>
- Dogbé, C. (2010). Modeling Crowd Dynamics by the Mean-Field Limit Approach. *Mathematical and Computer Modelling*, 52(9), 1506–1520. <https://doi.org/10.1016/j.mcm.2010.06.012>
- Elamvazhuthi, K., & Berman, S. (2019). Mean-Field Models in Swarm Robotics: A Survey. *Bioinspiration & Biomimetics*, 15(1), 015001. <https://doi.org/10.1088/1748-3190/ab49a4>
- Elamvazhuthi, K., Kakish, Z., Shirsat, A., & Berman, S. (2021). Controllability and Stabilization for Herding a Robotic Swarm Using a Leader: A Mean-Field Approach. *IEEE Transactions on Robotics*, 37(2), 418–432. <https://doi.org/10.1109/TRO.2020.3031237>
- Estrada-Rodriguez, G., & Gimpel, H. (2020). Interacting Particles with Lévy Strategies: Limits of Transport Equations for Swarm Robotic Systems. *SIAM Journal on Applied Mathematics*, 80(1), 476–498. <https://doi.org/10.1137/18M1205327>
- Hazra, G., Nandy, D., Kitchatinov, L., & Choudhuri, A. R. (2023). Mean Field Models of Flux Transport Dynamo and Meridional Circulation in the Sun and Stars. *Space Science Reviews*, 219(5). <https://doi.org/10.1007/s11214-023-00982-y>
- Huang, H., & Liu, J.-G. (2020). On the Mean-Field Limit for the Vlasov–Poisson–Fokker–Planck System. *Journal of Statistical Physics*, 25(12), 1915–1965. <https://doi.org/10.1007/s10955-020-02648-3>

- Huang, M., Malhamé, R., & Caines, P. (2006). Large population stochastic dynamic games: Closed-loop mckean-vlasov systems and the nash certainty equivalence principle. *Communications in Information and Systems*, 6(3), 221–252. <https://doi.org/10.4310/CIS.2006.v6.n3.a5>
- Karafyllis, I., Jiang, Z.-P., & Athanasiou, G. (2010). Nash equilibrium and robust stability in dynamic games: A small-gain perspective. *Computers & Mathematics with Applications*, 60(11), 2936–2952. <https://doi.org/10.1016/j.camwa.2010.09.054>
- Lasry, J. M., & Lions, P.-L. (2007). Mean Field Games. *Japanese Journal of Mathematics*, 2, 229–260. <https://doi.org/10.1007/s11537-007-0657-8>
- Ménec, S. L. (2024). Swarm guidance based on mean field game concepts. *International Game Theory Review*, 26(2), 244008. <https://doi.org/10.1142/S0219198924400085>
- Ornia, D. J., Zufiria, P. J., & Mazo, M. Jr. (2022). Mean Field Behavior of Collaborative Multiagent Foragers. *IEEE Transactions on Robotics*, 38(4), 2151–2165. <https://doi.org/10.1109/TRO.2022.3152691>
- Passos, Y. T., Duquesne, X., & Marcolino, L. S. (2023). Congestion control algorithms for robotic swarms with a common target based on the throughput of the target area. *Robotics and Autonomous Systems*, 159, 104284. <https://doi.org/10.1016/j.robot.2022.104284>
- Thouless, D. J., Anderson, P. W., & Palmer, R. G. (1977). Solution of 'Solvable Model of a Spin Glass'. *The Philosophical Magazine: A Journal of Theoretical Experimental and Applied Physics*, 35(3), 593–601. <https://doi.org/10.1080/14786437708235992>
- Vicsek, T., Czirók, A., Ben-Jacob, E., Cohen, I., & Shochet, O. (1995). Novel Type of Phase Transition in a System of Self-Driven Particles. *Physical Review Letters*, 75(6), 1226–1229. <https://doi.org/10.1103/PhysRevLett.75.1226>
- Yang, Y., Luo, R., Li, M., Zhou, M., Zhang, W., & Wang, J. (2018). Mean Field Multi-Agent Reinforcement Learning. In J. Dy & A. Krause (Eds.) *Proceedings of the 35th International Conference on Machine Learning*, (pp. 5571–5580, Vol. 80). PMLR.
- Zheng, T., Han, Q., & Lin, H. (2022). Transporting Robotic Swarms via Mean-Field Feedback Control. *IEEE Transactions on Automatic Control*, 67(8), 4170–4177. <https://doi.org/10.1109/TAC.2021.3108672>