Neural Mean-Field Games: Extending Mean-Field Game Theory with Neural Stochastic Differential Equations

Extended Abstract

Anna C.M. Thöni
Department of Artificial Intelligence,
Radboud University
Nijmegen, The Netherlands
chiara.thoeni@donders.ru.nl

Yoram Bachrach
Meta FAIR
London, United Kingdom
yorambac@meta.com

Tal Kachman
Department of Artificial Intelligence,
Radboud University
Nijmegen, The Netherlands
tal.kachman@donders.ru.nl

ABSTRACT

Mean-field game theory relies on approximating games that are intractible to model due to a very large, often infinite population of players. While these kinds of games can be solved analytically via the associated system of partial derivatives, this approach is not model-free, can lead to the loss of the existence or uniqueness of solutions, and may suffer from modelling bias. To reduce the dependency between the game and model, we introduce neural mean-field games: a combination of mean-field game theory and deep learning in the form of neural stochastic differential equations. The resulting model is data-driven, lightweight, and can learn extensive strategic interactions that are hard to capture using mean-field theory alone. In addition, the model is based on automatic differentiation, making it more robust and objective than approaches based on finite differences. We highlight the efficiency, robustness and flexibility of our approach by simulating viral dynamics based on real-world data.

KEYWORDS

Mean-field game theory; Neural stochastic differential equations; Nash Equilibria; Compartmental models

ACM Reference Format:
Anna C.M. Thöni, Yoram Bachrach, and Tal Kachman. 2026. Neural Mean-Field Games: Extending Mean-Field Game Theory with Neural Stochastic Differential Equations: Extended Abstract. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 3 pages. https://doi.org/10.65109/WCLX3778

1 INTRODUCTION

Game theory investigates strategic interactions between self-interested agents that inhabit the same environment and are affected by each other's actions. Such analysis has applications in economics [3, 15, 17], crowd motions [2, 8] and evolutionary biology [4]. For games with large populations, it can be intractable to explicitly model all interactions between agents. Instead, such games are approximated using mean-field game (MFG) theory [19, 21]. MFG theory replaces the agents with their mean field, effectively summarizing their interactions by a distribution that is described by two coupled partial differential equations (PDEs) [18]. Even though the population dynamics are subject to the actions of the individual agents, the agents are too small to exert an independent influence. Instead, all agents are expected to jointly converge towards an optimal control, arriving at a state that can be considered the game's Nash Equilibrium (NE). While the NE can be found by numerically solving the PDEs that accompany the game [10] using finite differences [1, 9], such a discretization introduces errors and stability issues that are subject to the chosen method and may ultimately lead to the loss of the existence or uniqueness of the solutions [36]. Instead, the solutions to the mean-field game can be approximated using artificial neural networks [23], in the form of recurrent architectures [11, 12], deep policy iteration [7], deep Galerkin methods [6], or actor-critic-based approaches [16, 35, 36]. Yet, all approaches mentioned above require Fictitious Play [5, 22, 27, 28, 31] or full knowledge of the game of interest, making the solution either numerically inexact, limiting the models' applications, or introducing a modelling bias.

Our Contribution: To retain the numerical exactness of the differential equation-based solutions without requiring full knowledge of the game of interest, we analyze MFGs using neural stochastic differential equations (SDEs; [24, 34]). Neural SDEs form the stochastic extension of neural ordinary differential equations (ODEs), a family of neural networks introduced by Chen et al. [13]. To model the MFG dynamics as a neural SDE, we complement the existing system of coupled differential equations with a neural network.

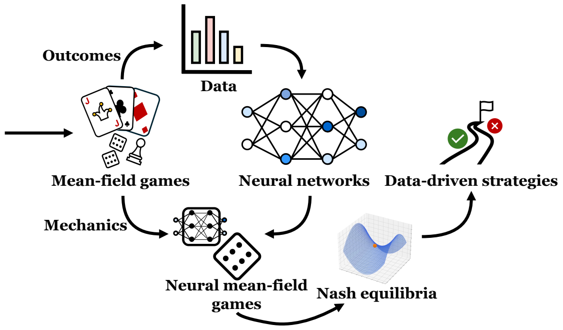

Figure 1: An overview of modelling MFGs with neural SDEs. The resulting neural mean-field game combines the MFG mechanics with the neural network output to create more informed strategies.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/WCLX3778

The resulting model, shown in Figure 1, is data-driven, model-free, and can be solved using automatic differentiation. We demonstrate the flexibility of the neural MFGs by modelling the dynamics of COVID-19 infections in Japan. Dynamical systems for epidemics are highly nonlinear and depend on various features, including population parameters, restrictive measures, and viral characteristics that are often beyond the scope of classical mathematical models [25]. To provide insight into the complex mechanics behind epidemic systems, we extend the classical models by learning viral mechanics from a large range of epidemic characteristics.

Data, Code and Paper Availability The work uses data obtained from The Japanese Ministry of Health, Labour and Welfare [32] and Mathieu et al. [26]. The code associated with the model implementation and analyses is publicly available on GitHub. For full algorithmic details, results, and related work, we refer the reader to the full paper version [33].

2 EXPERIMENTS AND RESULTS

We model viral dynamics as an MFG with a finite number of states and consider a population of homogeneous players in continuous time $t \in [0, T]$. We consider three states in total, where each player can be susceptible (S), infected (I), or recovered (R). The mean-field of the game, $m(t)$, summarizes the fraction of players in each state. The resulting SIR model forms a Markov chain with non-linear dynamics: the transition rates do not merely depend on the players' actions, but are influenced by their interaction with the population distribution m as well [14, 20]. In addition, we incorporate auxiliary measures influencing state transition rates, including vaccinations, school closings, lockdowns, transport restrictions and the 2020 Summer Olympics held in Tokyo from 23 July until 8 August 2021. We include these observations in the neural MFG by encoding them as a one-hot vector and actively learning their influence on the virus dynamics from the data.

The predictions made with the neural MFG, shown in Figure 2a, capture the peaks in infections around January, May, and August 2021 (middle panel). As a baseline, we compare the Neural MFG predictions to a similar model without a neural drift component: here, we model the SIR dynamics using predetermined, deterministic drift-based, and a neural diffusion with the same parameterization as the diffusion in the Neural MFG. The predictions of the resulting model, presented in Figure 2b, are less accurate than those from the neural MFG. Even though the predictions capture the shape of the viral dynamics, their scale is incorrect. In addition, the peak of infections observed around May 2021 is overestimated, suggesting that the addition of auxiliary variables enables the neural MFG to more accurately predict the viral dynamics.

3 CONCLUSION

Mean-field game theory is used to efficiently model games with a very large population of agents. The agents' Nash equilibria are commonly found via reinforcement learning or by solving the system of PDEs associated with the MFG. Both methods have their advantages, the first being model-free and the second numerically exact. To get the best of both worlds, we extend the PDE formulation of the MFG with a neural network. The resulting neural MFG exerts a data-driven influence on the game's Nash equilibria, deriving strategies that are in line with observed game outcomes, auxiliary metrics, or past dynamics. Leveraging those properties, the model can fill in dynamics that are unaccounted for in the original MFG specification. In addition, neural MFGs are a numerically exact and lightweight alternative to (deep) RL-based methods for solving MFGs. Here, their dependency on automatic differentiation makes the neural MFGs more objective than approaches that solve PDEs using finite differences [13, 30], which may prevent the loss of existence or uniqueness of solutions associated with solving MFGs [36]. We translate the theoretical advantages stated above into an impactful application by successfully modelling the outbreak of an infectious disease. Hereby, our work helps bridge the gap between mathematical modelling and real-world applications. Even though we highlight a biological application, we expect the model to find applications in any system associated with non-atomic, anonymous players and observations. Such systems are, e.g., also commonly found in markets and economics [29].

ACKNOWLEDGMENTS

Tal Kachman and Anna C.M. Thöni acknowledge funding from the National Growth Fund project "Big Chemistry" (1420578) and the European Laboratory for Learning and Intelligent Systems (ELLIS). The authors also thank Mathieu Laurière for the insightful discussions.

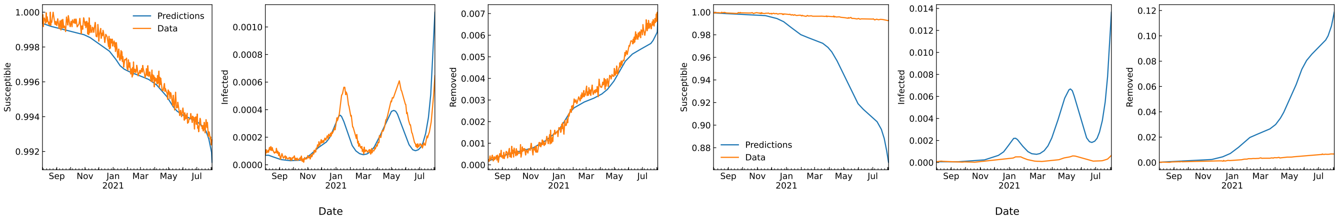

(a)

(b)

Figure 2: Comparison of the noisy, observed data (orange) and the predictions (blue) made with (a) the neural MFG and (b) a neural SDE with a predefined, deterministic drift for the proportions of the Japanese population that are susceptible, infected, or removed. The results are normalized with respect to the total population to the domain [0, 1] and consider the period between October 1st, 2020, and October 3rd, 2021.

REFERENCES

[1] Yves Achdou and Italo Capuzzo-Dolcetta. 2010. Mean Field Games: Numerical Methods. *SIAM J. Numer. Anal.* 48, 3 (2010), 1136–1162. https://doi.org/10.1137/090758477

[2] Yves Achdou, Pierre Cardaliaguet, François Delarue, Alessio Porretta, and Filippo Santambrogio. 2019. *Mean Field Games.* Vol. 2281. Springer, Heidelberg, Germany. https://doi.org/10.1007/978-3-030-59837-2

[3] Yves Achdou, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. 2021. Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach. *The Review of Economic Studies* 89, 1 (04 2021), 45–86. https://doi.org/10.1093/restud/rdab002

[4] J. McKenzie Alexander. 2021. Evolutionary Game Theory. In *The Stanford Encyclopedia of Philosophy* (Summer 2021 ed.), Edward N. Zalta (Ed.). Metaphysics Research Lab, Stanford University, Stanford, CA. https://doi.org/10.1017/9781108582063

[5] Berkay Anahtarci, Can Deha Kariksiz, and Naci Saldi. 2023. Q-Learning in Regularized Mean-field Games. *Dynamic Games and Applications* 13, 1 (2023), 89–117. https://doi.org/10.1007/s13235-022-00450-2

[6] Mouhcine Assouli and Badr Missaoui. 2023. Deep learning for Mean Field Games with non-separable Hamiltonians. *Chaos, Solitons & Fractals* 174 (2023), 113802. https://doi.org/10.1016/j.chaos.2023.113802

[7] Mouhcine Assouli and Badr Missaoui. 2024. Deep Policy Iteration for high-dimensional mean field games. *Appl. Math. Comput.* 481 (2024), 128923. https://doi.org/10.1016/j.amc.2024.128923

[8] Alexander Aurell and Boualem Djehiche. 2019. Modeling tagged pedestrian motion: A mean-field type game approach. *Transportation Research Part B: Methodological* 121 (2019), 168–183. https://doi.org/10.1016/j.trb.2019.01.011

[9] Luis Briceño-Arias, Dante Kalise, Ziad Kobeissi, Mathieu Laurière, Álvaro Mateos González, and Francisco Silva. 2019. On the implementation of a primal-dual algorithm for second order time-dependent Mean Field Games with local couplings. *ESAIM: Proceedings and Surveys* 65 (01 2019), 330–348. https://doi.org/10.1051/proc/201965330

[10] Pierre Cardaliaguet. 2013. Notes on Mean Field Games.

[11] René Carmona and Mathieu Laurière. 2021. Convergence Analysis of Machine Learning Algorithms for the Numerical Solution of Mean Field Control and Games I: The Ergodic Case. *SIAM J. Numer. Anal.* 59, 3 (2021), 1455–1485. https://doi.org/10.1137/19M1274377

[12] Rene Carmona and Mathieu Laurière. 2022. Convergence analysis of machine learning algorithms for the numerical solution of mean field control and games: II—the finite horizon case. *The Annals of Applied Probability* 32 (12 2022). https://doi.org/10.1214/21-AAP1715

[13] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K. Duvenaud. 2018. Neural Ordinary Differential Equations. *Advances in Neural Information Processing Systems* 31 (2018). https://doi.org/10.48550/arXiv.1806.07366

[14] Josu Doncel, Nicolas Gast, and Bruno Gaujal. 2016. Mean-Field Games with Explicit Interactions. *Preprint hal-01277098* (2016). https://inria.hal.science/hal-01277098v1 working paper or preprint.

[15] David Easley, Jon Kleinberg, et al. 2010. *Networks, crowds, and markets: Reasoning about a highly connected world.* Vol. 1. Cambridge university press, Cambridge, UK. https://doi.org/10.1017/CBO9780511761942

[16] Noufel Frikha, Maximilien Germain, Mathieu Laurière, Huyên Pham, and Xuanye Song. 2025. Actor-Critic learning for mean-field control in continuous time. *Journal of Machine Learning Research* 26, 127 (2025), 1–42.

[17] Robert Gibbons. 1992. *Game Theory for Applied Economists.* Princeton University Press, Princeton, NJ.

[18] Olivier Guéant, Jean-Michel Lasry, and Pierre-Louis Lions. 2011. *Mean Field Games and Applications.* Springer Berlin Heidelberg, Berlin, Heidelberg, 205–266. https://doi.org/10.1007/978-3-642-14660-2_3

[19] Minyi Huang, Roland P. Malhamé, and Peter E. Caines. 2006. Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle. *Communications in Information & Systems* 6, 3 (2006), 221 – 252. https://doi.org/10.4310/CIS.2006.v6.n3.a5

[20] Vassili N. Kolokoltsov. 2010. *Nonlinear Markov Processes and Kinetic Equations.* Cambridge University Press, Cambridge. https://doi.org/10.1017/CBO9780511760303

[21] Jean-Michel Lasry and Pierre-Louis Lions. 2007. Mean Field Games. *Japanese Journal of Mathematics* 2 (03 2007), 229–260. https://doi.org/10.1007/s11537-007-0657-8

[22] Timothy Lillicrap, Jonathan Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. 2015. Continuous Control with Deep Reinforcement Learning. *CoRR* (09 2015). https://doi.org/10.48550/arXiv.1509.02971

[23] Alex Tong Lin, Samy Wu Fung, Wuchen Li, Levon Nurbekyan, and Stanley J. Osher. 2021. Alternating the population and control neural networks to solve high-dimensional stochastic mean-field games. *Proceedings of the National Academy of Sciences* 118, 31 (2021), e2024713118. https://doi.org/10.1073/pnas.2024713118

[24] Xuanqing Liu, Tesi Xiao, Si Si, Qin Cao, Sanjiv Kumar, and Cho-Jui Hsieh. 2019. Neural SDE: Stabilizing Neural ODE Networks with Stochastic Noise. arXiv:1906.02355 [cs.LG] https://arxiv.org/abs/1906.02355

[25] Tchavdar T. Marinov and Rossitza S. Marinova. 2022. Adaptive SIR model with vaccination: simultaneous identification of rates and functions illustrated with COVID-19. *Scientific Reports* 12, 1 (2022), 15688. https://doi.org/10.1038/s41598-022-20276-7

[26] Edouard Mathieu, Hannah Ritchie, Lucas Rodés-Guirao, Cameron Appel, Daniel Gavrilov, Charlie Giattino, Joe Hasell, Bobbie Macdonald, Saloni Dattani, Diana Beltekian, Esteban Ortiz-Ospina, and Max Roser. 2020. COVID-19 Pandemic. *Our World in Data* (2020). https://ourworldindata.org/coronavirus.

[27] Rajesh Mishra, Deepanshu Vasal, and Sriram Vishwanath. 2023. Model-Free Reinforcement Learning for Mean Field Games. *IEEE Transactions on Control of Network Systems* 10, 4 (2023), 2141–2151. https://doi.org/10.1109/TCNS.2023.3264934

[28] Sarah Perrin, Mathieu Laurière, Pérolat Julien, Matthieu Geist, Romuald Elie, and Olivier Pietquin. 2021. Mean Field Games Flock! The Reinforcement Learning Way. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence (IJCAI-21).* 356–362. https://doi.org/10.24963/ijcai.2021/50

[29] David Schmeidler. 1973. Equilibrium Points of Non-Atomic Games. *Journal of Statistical Physics* 7 (May 1973), 295–300. https://doi.org/10.1007/BF01014905

[30] J. Su and J. E. Renaud. 1997. Automatic Differentiation in Robust Optimization. *AIAA Journal* 35, 6 (June 1997), 1072–1079. https://doi.org/10.2514/2.196

[31] Sriram Ganapathi Subramanian, Matthew E. Taylor, Mark Crowley, and Pascal Poupart. 2021. Partially Observable Mean Field Reinforcement Learning. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems.* 537–545. https://doi.org/10.48550/arXiv.2012.15791

[32] The Japanese Ministry of Health, Labour and Welfare. [n.d.]. Regarding the number of vaccinations of the new coronavirus vaccine during the special temporary vaccination period. https://www.mhlw.go.jp/stf/seisakunitsuite/bunya/kenkou_iryou/kenkou/kekkaku-kansenshou/yobou-sesshu/syukeihou_00002.html. Accessed on 12-08-2025.

[33] Anna C. M. Thöni, Yoram Bachrach, and Tal Kachman. 2025. Neural Mean-Field Games: Extending Mean-Field Game Theory with Neural Stochastic Differential Equations. https://doi.org/10.48550/arXiv.2504.13228 arXiv:2504.13228 [cs.LG]

[34] Belinda Tzen and Maxim Raginsky. 2019. Neural Stochastic Differential Equations: Deep Latent Gaussian Models in the Diffusion Limit. arXiv:1905.09883 [cs.LG] https://arxiv.org/abs/1905.09883

[35] Zejian Zhou and M. Sami Fadali and Hao Xu. 2020. Biomimetic Optimal Tracking Control using Mean Field Games and Spiking Neural Networks. *IFAC-PapersOnLine* 53, 2 (2020), 8112–8117. https://doi.org/10.1016/j.ifacol.2020.12.2281 21st IFAC World Congress.

[36] Zejian Zhou, Yuzhu Zhang, and Hao Xu. 2021. Reinforcement Learning-based Decentralized Optimal Control for Large-Scale Multi-agent System by Using Neural Networks and Discrete-time Mean Field Games. In *2021 International Joint Conference on Neural Networks (IJCNN).* 1–6. https://doi.org/10.1109/IJCNN52387.2021.9534288