Proceedings of Machine Learning Research vol 247:1–6, 2024 37th Annual Conference on Learning Theory

Open problem: Convergence of single-timescale mean-field Langevin descent-ascent for two-player zero-sum games

Guillaume Wang GUILLAUME.WANG@EPFL.CH and **Lénaïc Chizat** LENAIC.CHIZAT@EPFL.CH
Institute of Mathematics, École polytechnique fédérale de Lausanne (EPFL), Station Z, CH-1015 Lausanne

Editors: Shipra Agrawal and Aaron Roth

Abstract

Let a smooth function $f : \mathbb{T}^d \times \mathbb{T}^d \to \mathbb{R}$ over the d-torus and $\beta > 0$. Consider the min-max objective functional $F_\beta(\mu,\nu) = \iint f \mathrm{d}\mu\mathrm{d}\nu + \beta^{-1}H(\mu) - \beta^{-1}H(\nu)$ over $\mathcal{P}(\mathbb{T}^d) \times \mathcal{P}(\mathbb{T}^d)$, where H denotes the negative differential entropy. Its unique saddle point defines the entropy-regularized mixed Nash equilibrium of a two-player zero-sum game, and its Wasserstein gradient descent-ascent flow (μ_t, ν_t) corresponds to the mean-field limit of a Langevin descent-ascent dynamics. Do μ_t and ν_t converge (weakly, say) as $t \to \infty$, for any f and β? This rather natural qualitative question is still open, and it is not clear whether it can be addressed using the tools currently available for the analysis of dynamics in Wasserstein space. Even though the simple trick of using a different timescale for the ascent versus the descent is known to guarantee convergence, we propose this question as a toy setting to further our understanding of the Wasserstein geometry for optimization.

Keywords: Wasserstein gradient flow, mean-field Langevin dynamics, min-max optimization

1. Introduction

The use of the minimax two-player game framework in Generative Adversarial Networks (GANs) (Goodfellow et al., 2014) and distributionally robust learning (Sinha et al., 2018; Madry et al., 2018) inspired renewed interest in the question of algorithmically identifying Nash equilibria of two-player zero-sum games, or equivalently, saddle points of min-max problems $\min_{x\in\mathcal{X}} \max_{y\in\mathcal{Y}} f(x,y)$, given first-order access to f. In order to bypass issues of existence and uniqueness of solutions, one line of work has focused on finding (entropy-regularized) mixed Nash equilibria (Hsieh et al., 2019), defined as follows.

Definition 1 *Let $\mathcal{X} = \mathcal{Y} = \mathbb{T}^d$ the Euclidean torus and $f : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$ a C^2 function.*[1] *A mixed-strategy Nash equilibrium (MNE) is a couple of probability measures $(\mu^*, \nu^*) \in \mathcal{P}(\mathcal{X}) \times \mathcal{P}(\mathcal{Y})$ such that (all double integrals are over $\mathcal{X} \times \mathcal{Y}$)*

$$\forall(\mu,\nu) \in \mathcal{P}(\mathcal{X}) \times \mathcal{P}(\mathcal{Y}), \quad \iint f \, \mathrm{d}\mu^* d\nu \le \iint f \, \mathrm{d}\mu^*\mathrm{d}\nu^* \le \iint f \, \mathrm{d}\mu\mathrm{d}\nu^*.$$

Observe that, equivalently, a MNE is a saddle point of the infinite-dimensional min-max problem $\min_{\mu\in\mathcal{P}(\mathcal{X})} \max_{\nu\in\mathcal{P}(\mathcal{Y})} \iint f\mathrm{d}\mu\mathrm{d}\nu$. For $\beta > 0$, we call entropy-regularized MNE the saddle point of

$$\min_{\mu\in\mathcal{P}(\mathcal{X})} \max_{\nu\in\mathcal{P}(\mathcal{Y})} F_\beta(\mu,\nu), \qquad F_\beta(\mu,\nu) = \iint f \, \mathrm{d}\mu\mathrm{d}\nu + \beta^{-1}H(\mu) - \beta^{-1}H(\nu), \tag{1}$$

where $H(\mu) = \int_{\mathcal{X}} \log\left(\frac{\mathrm{d}\mu}{\mathrm{d}x}\right) \mathrm{d}\mu$ is the (negative) differential entropy.

1. We purposefully present the problem in the simplest setting. Possible extensions include taking as $\mathcal{X}, \mathcal{Y}$ compact Riemannian manifolds without boundaries (Domingo-Enrich et al., 2020), or $\mathcal{X}, \mathcal{Y} = \mathbb{R}^d$ with additional assumptions on f and additional confining terms in (2) (Kim et al., 2024).

The availability of first-order access to f prompted the use of Wasserstein gradient descent-ascent flows (WGFs) for this setting (Domingo-Enrich et al., 2020), instead of multiplicative-weight methods as more traditionally considered for games with finite strategy sets $\mathcal{X}, \mathcal{Y}$ (Wei et al., 2020; Cen et al., 2021) (which only use zeroth-order access to f). The WGF of F_β is given by the PDE

$$\begin{cases} \partial_t \mu_t = \ \ \operatorname{div}\left(\mu_t \nabla_x \left[\int_{\mathcal{Y}} f(x,y)\mathrm{d}\nu_t(y)\right]\right) + \beta^{-1}\Delta\mu_t \ \text{ over } \mathcal{X} \\ \partial_t \nu_t = -\operatorname{div}\left(\nu_t \nabla_y \left[\int_{\mathcal{X}} f(x,y)\mathrm{d}\mu_t(x)\right]\right) + \beta^{-1}\Delta\nu_t \quad \text{over } \mathcal{Y}, \end{cases} \tag{2}$$

which can be viewed as the mean-field limit ($N \to \infty$) of the system of SDEs, called Langevin descent-ascent dynamics,

$$\begin{cases} \forall i \leq N, \ \mathrm{d}X_t^i = -\frac{1}{N}\sum_{j=1}^N \nabla_x f(X_t^i, Y_t^j)\mathrm{d}t + \sqrt{2\beta^{-1}}\mathrm{d}B_t^i \\ \forall j \leq N, \ \mathrm{d}Y_t^j = \frac{1}{N}\sum_{i=1}^N \nabla_y f(X_t^i, Y_t^j)\mathrm{d}t + \sqrt{2\beta^{-1}}\mathrm{d}B_t^j \end{cases} \tag{3}$$

via $\mu_t = \frac{1}{N}\sum_{i=1}^N \delta_{X_t^i}$ and $\nu_t = \frac{1}{N}\sum_{j=1}^N \delta_{Y_t^j}$ (Domingo-Enrich et al., 2020, Thm. 3 (i)). For this reason, we follow Kim et al. (2024) in referring to (2) as *mean-field Langevin descent-ascent* (MFL-DA).

We pose the open question:

Does the MFL-DA trajectory (μ_t, ν_t) converge weakly to some (μ^*, ν^*) as $t \to \infty$?

It is proved in Domingo-Enrich et al. (2020, Thm. 1) that if (μ_t, ν_t) converges weakly, then the limit (μ^*, ν^*) is the unique saddle point of F_β. However, despite several positive results concerning variants of MFL-DA (reviewed hereafter), the convergence of Eq. (2) itself has remained elusive.

1.1. Other notions of convergence

The natural measure of suboptimality for the min-max problem $\min_{\mathcal{P}(\mathcal{X})} \max_{\mathcal{P}(\mathcal{Y})} F_\beta$ is the Nikaido-Isoda (NI) error defined by

$$\mathrm{NI}(\mu,\nu) = \max_{\nu' \in \mathcal{P}(\mathcal{Y})} F_\beta(\mu,\nu') - \min_{\mu' \in \mathcal{P}(\mathcal{X})} F_\beta(\mu',\nu).$$

One can show that there exists a constant $\rho_\beta^* > 0$ such that[2]

$$\frac{\rho_\beta^*}{2}\left(W_2^2(\mu,\mu^*) + W_2^2(\nu,\nu^*)\right) \leq \mathrm{KL}\left(\mu\|\mu^*\right) + \mathrm{KL}\left(\nu\|\nu^*\right) \leq \beta\, \mathrm{NI}(\mu,\nu) \tag{4}$$

where W_2 is the 2-Wasserstein distance and $\mathrm{KL}\left(\cdot\|\cdot\right)$ is the relative entropy. In particular, the following variants of the open question are increasingly stronger: if (μ_t, ν_t) is the MFL-DA trajectory,

- *Does it hold $W_2^2(\mu_t,\mu^*) + W_2^2(\nu_t,\nu^*) \to 0$ as $t \to \infty$?* Since the Wasserstein distance metrizes weak convergence, this is equivalent to our main open question.
- Convergence in relative entropy: *Does it hold* $\mathrm{KL}\left(\mu_t\|\mu^*\right) + \mathrm{KL}\left(\nu_t\|\nu^*\right) \to 0$*?*
- Convergence in NI error: *Does it hold* $\mathrm{NI}(\mu_t,\nu_t) \to 0$*?*

2. The first inequality of (4) follows by noting that μ^* (resp. ν^*) satisfies a logarithmic Sobolev inequality (LSI) (Lu, 2023, Lemma 2.1), and so a Talagrand (T2) inequality (Otto and Villani, 2000) with the same constant. The second inequality is proved e.g. in Kim et al. (2024, Lemma 3.5).

Yet other variants of the open question can be formulated by asking only for local convergence, i.e., by assuming that the initialization (μ_0, ν_0) lies in some W_2 or relative entropy neighborhood of (μ^*, ν^*), or in some sublevel set of NI. Once the question of qualitative convergence is settled, a natural development may be to ask for explicit rates of convergence.

1.2. Motivation

This open question touches upon three different topics. In terms of **design and analysis of algorithms**, the question is about the convergence of a particular continuous optimization dynamics, MFL-DA, for the min-max problem (1). In fact, it is known that convergence can be guaranteed for certain modifications of MFL-DA (Lu, 2023; Kim et al., 2024) (reviewed below). The question is whether these modifications are really necessary for convergence to hold.

The open question is also interesting in its own right from the perspective of **game theory** and economics, where it essentially asks to characterize the long-time behavior of a system of $2N$ agents interacting via the SDE system (3), in the mean-field limit $N \gg 1$.

From a broader point of view, the open question is about understanding the use of the **Wasserstein geometry for optimization**. A number of common machine learning tasks can be framed as optimization over the space of probability measures via noisy particle methods (of which (3) is an example), including sampling (Wibisono, 2018), training of two-layer neural networks (Mei et al., 2018), trajectory inference (Chizat et al., 2022), quantization of measures (Xu et al., 2022), or grid-free regularized Wasserstein barycenters (Chizat, 2023). Theoretical analyses for those tasks thus borrow from and extend our understanding of Wasserstein geometry for optimization. This topic has steadily progressed in the recent years and now relies on solid foundations, although several important questions remain open. Yet, MFL-DA is a rare instance of a Wasserstein optimization dynamics for which *even qualitative convergence guarantees are not known*. Regardless of whether the answer to the open question turns out to be positive or negative, there is hope that progress for this toy setting will lead to a better understanding of WGFs as a whole.

2. Related work and state of the problem

Positive results for variants of MFL-DA. The works of Mei et al. (2018); Hu et al. (2021); Chizat (2022); Nitanda et al. (2022) analyze the convergence of mean-field Langevin descent dynamics, i.e., the first half of Eq. (2) where $\nabla_x \big[\int_{\mathcal{Y}} f(x,y) \mathrm{d}\nu_t(y) \big]$ is replaced by the Wasserstein gradient at μ_t of a convex functional $G : \mathcal{P}(\mathcal{X}) \to \mathbb{R}$. It is a natural analog of MFL-DA for minimization instead of min-max.

A two-timescale variant of (2), where the right-hand side of the equation for $\partial_t \nu_t$ is multiplied by a small fixed $\varepsilon > 0$, has been studied by Lu (2023) (and previously by Ma and Ying (2021) for infinitesimal ε). It is shown that it is convergent in NI error for any $\varepsilon \leq \varepsilon_0$, for some constant ε_0 dependent on f and β. The main proof ingredient is that for any ν, the "min" objective $F_\beta(\cdot, \nu) : \mathcal{P}(\mathcal{X}) \to \mathbb{R}$ satisfies a Wasserstein-space analog of the Polyak-Lojasiewicz (PL) inequality with a constant $\beta^{-1}\rho_\beta$ independent of ν, and symmetrically for the "max" objective; in other words, F_β satisfies an analog of the *two-sided PL inequality condition* (Yang et al., 2020). This allows to mimick the convergence analysis of two-timescale gradient descent-ascent flow in finite dimension under this condition (Doan, 2022). Note that there exist finite-dimensional min-max problems satisfying this condition and for which, numerically, single-timescale gradient flow seems not to converge (Yang et al., 2020, Remark 1).

A "time-averaged gradient" variant of (2) has been studied by Kim et al. (2024), where the Wasserstein gradient $\nabla_x \left[\int_{\mathcal{Y}} f(x,y) \mathrm{d}\nu_t(y) \right]$ is replaced by $\int_0^t \alpha_s \nabla_x \left[\int_{\mathcal{Y}} f(x,y) \mathrm{d}\nu_s(y) \right] \mathrm{d}s / \left(\int_0^t \alpha_s \mathrm{d}s \right)$ for some weighting $\alpha : \mathbb{R}_+ \to \mathbb{R}_+$, and likewise for $\nabla_y \left[\int_{\mathcal{X}} f(x,y) d\mu_t(y) \right]$. It is shown that for $\alpha_t = t^r$ for a fixed $r > -1$, the dynamics converges in NI error (by specializing their Prop. 3.3 to $\mathcal{L}$ being bilinear).

Positive results in specific cases. As mentioned in Domingo-Enrich et al. (2020, Sec. 4.1), the high temperature case $\beta^{-1} \gg 1$ can be analyzed using generic tools developed for diffusion or McKean-Vlasov processes (Eberle et al., 2019). More precisely, one can show that MFL-DA is convergent in NI error for $\beta^{-1} \gtrsim \max \left(\|\nabla_x f\|_\infty, \|\nabla_y f\|_\infty \right)$, or for $\beta^{-1} \gtrsim \max \left(\|\nabla_x \nabla_y\|_\infty, \|f\|_{\mathrm{osc}} \right\}$ where $\|f\|_{\mathrm{osc}} = \sup f - \inf f$, where "$\gtrsim$" hides universal constants when $\mathcal{X} = \mathcal{Y} = \mathbb{T}^d$ (or more generally, constants dependent only on $\mathcal{X}, \mathcal{Y}$).

Other relevant results. Other relevant results are contained in Sec. 4, 5 of the retracted arXiv paper Domingo-Enrich and Bruna (2022) (and those sections are not affected by the mistake that caused the retraction). Its Sec. 5 shows that if $\mathcal{X} = \mathcal{Y} = \mathbb{R}$ and $f(x, \cdot)$ and $f(\cdot, y)$ are quadratic for each x, y, then for μ_0 and ν_0 being Gaussians located close to μ^*, ν^* in a certain sense, we have $(\mu_t, \nu_t) \to (\mu^*, \nu^*)$ weakly. Its Sec. 4 shows an example of a function $G_\beta : \mathcal{P}(\mathbb{R}) \times \mathcal{P}(\mathbb{R}) \to \mathbb{R}$ which is convex-concave and displacement-convex-concave, whose WGF exhibits a cycling behavior. Although these results do not quite fit the setting considered here as $\mathcal{X}$ and $\mathcal{Y}$ are non-compact, the first one suggests that local convergence of MFL-DA may hold, and the second one may offer a path to constructing counter-examples.

References

Shicong Cen, Yuting Wei, and Yuejie Chi. Fast policy extragradient methods for competitive games with entropy regularization. *Advances in Neural Information Processing Systems*, 34:27952–27964, 2021.

Lénaïc Chizat. Doubly regularized entropic wasserstein barycenters. *arXiv preprint arXiv:2303.11844*, 2023.

Lénaïc Chizat. Mean-field langevin dynamics: Exponential convergence and annealing. *Transactions on Machine Learning Research*, 2022.

Lénaïc Chizat, Stephen Zhang, Matthieu Heitz, and Geoffrey Schiebinger. Trajectory inference via mean-field langevin in path space. *Advances in Neural Information Processing Systems*, 35: 16731–16742, 2022.

Thinh Doan. Convergence rates of two-time-scale gradient descent-ascent dynamics for solving nonconvex min-max problems. In *Learning for Dynamics and Control Conference*, pages 192–206. PMLR, 2022.

Carles Domingo-Enrich and Joan Bruna. Simultaneous transport evolution for minimax equilibria on measures. *arXiv preprint arXiv:2202.06460*, 2022.

Carles Domingo-Enrich, Samy Jelassi, Arthur Mensch, Grant Rotskoff, and Joan Bruna. A mean-field analysis of two-player zero-sum games. *Advances in neural information processing systems*, 33:20215–20226, 2020.

Andreas Eberle, Arnaud Guillin, and Raphael Zimmer. Quantitative harris-type theorems for diffusions and mckean–vlasov processes. *Transactions of the American Mathematical Society*, 371 (10):7135–7173, 2019.

Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.

Ya-Ping Hsieh, Chen Liu, and Volkan Cevher. Finding mixed nash equilibria of generative adversarial networks. In *International Conference on Machine Learning*, pages 2810–2819. PMLR, 2019.

Kaitong Hu, Zhenjie Ren, David Šiška, and Łukasz Szpruch. Mean-field langevin dynamics and energy landscape of neural networks. In *Annales de l'Institut Henri Poincare (B) Probabilites et statistiques*, volume 57, pages 2043–2065. Institut Henri Poincaré, 2021.

Juno Kim, Kakei Yamamoto, Kazusato Oko, Zhuoran Yang, and Taiji Suzuki. Symmetric mean-field langevin dynamics for distributional minimax problems. In *International Conference on Learning Representations*, 2024.

Yulong Lu. Two-scale gradient descent ascent dynamics finds mixed nash equilibria of continuous games: A mean-field perspective. In *International Conference on Machine Learning*, pages 22790–22811. PMLR, 2023.

Chao Ma and Lexing Ying. Provably convergent quasistatic dynamics for mean-field two-player zero-sum games. In *International Conference on Learning Representations*, 2021.

Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018.

Song Mei, Andrea Montanari, and Phan-Minh Nguyen. A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, 115(33):E7665–E7671, 2018.

Atsushi Nitanda, Denny Wu, and Taiji Suzuki. Convex analysis of the mean field langevin dynamics. In *International Conference on Artificial Intelligence and Statistics*, pages 9741–9757. PMLR, 2022.

Felix Otto and Cédric Villani. Generalization of an inequality by talagrand and links with the logarithmic sobolev inequality. *Journal of Functional Analysis*, 173(2):361–400, 2000.

Aman Sinha, Hongseok Namkoong, and John Duchi. Certifying some distributional robustness with principled adversarial training. In *International Conference on Learning Representations*, 2018.

Chen-Yu Wei, Chung-Wei Lee, Mengxiao Zhang, and Haipeng Luo. Linear last-iterate convergence in constrained saddle-point optimization. In *International Conference on Learning Representations*, 2020.

Andre Wibisono. Sampling as optimization in the space of measures: The langevin dynamics as a composite optimization problem. In *Conference on Learning Theory*, pages 2093–3027. PMLR, 2018.

Lantian Xu, Anna Korba, and Dejan Slepcev. Accurate quantization of measures via interacting particle-based optimization. In *International Conference on Machine Learning*, pages 24576–24595. PMLR, 2022.

Junchi Yang, Negar Kiyavash, and Niao He. Global convergence and variance reduction for a class of nonconvex-nonconcave minimax problems. *Advances in Neural Information Processing Systems*, 33:1153–1165, 2020.