

Concave Utility Reinforcement Learning: the Mean-Field Game Viewpoint

Matthieu Geist
Google Research, Brain Team

Julien Pérolat
Deepmind Paris

Mathieu Laurière
Google Research, Brain Team

Romuald Elie
Deepmind Paris

Sarah Perrin
Univ. Lille, CNRS, Inria, Centrale Lille,
UMR 9189 CRISTAL

Olivier Bachem
Google Research, Brain Team

Rémi Munos
Deepmind Paris

Olivier Pietquin
Google Research, Brain Team

ABSTRACT

Concave Utility Reinforcement Learning (CURL) extends RL from linear to concave utilities in the occupancy measure induced by the agent’s policy. This encompasses not only RL but also imitation learning and exploration, among others. Yet, this more general paradigm invalidates the classical Bellman equations, and calls for new algorithms. Mean-field Games (MFGs) are a continuous approximation of many-agent RL. They consider the limit case of a continuous distribution of identical agents, anonymous with symmetric interests, and reduce the problem to the study of a single representative agent in interaction with the full population. Our core contribution consists in showing that CURL is a subclass of MFGs. We think this important to bridge together both communities. It also allows to shed light on aspects of both fields: we show the equivalence between concavity in CURL and monotonicity in the associated MFG, between optimality conditions in CURL and Nash equilibrium in MFG, or that Fictitious Play (FP) for this class of MFGs is simply Frank-Wolfe, bringing the first convergence rate for discrete-time FP for MFGs. We also experimentally demonstrate that, using algorithms recently introduced for solving MFGs, we can address the CURL problem more efficiently.

KEYWORDS

Concave Utility Reinforcement Learning; Mean-field Games

ACM Reference Format:

Matthieu Geist, Julien Pérolat, Mathieu Laurière, Romuald Elie, Sarah Perrin, Olivier Bachem, Rémi Munos, and Olivier Pietquin. 2022. Concave Utility Reinforcement Learning: the Mean-Field Game Viewpoint. In *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022)*, Online, May 9–13, 2022, IFAAMAS, 9 pages.

1 INTRODUCTION

Reinforcement Learning (RL) aims at finding a policy maximizing the expected discounted cumulative reward along trajectories generated by the policy. This objective can be rewritten as an inner product between the occupancy measure induced by the policy (i.e., the discounted distribution of state-action pairs it visits) and a policy-independent reward for each state-action pair. Yet, several

recently investigated problems aim at optimizing a more general function of this occupancy measure: Hazan et al. [21] propose pure exploration in RL as maximizing the entropy of the policy-induced state occupancy measure, and Ghasemipour et al. [16] showed that many imitation learning (IL) algorithms minimize an f -divergence between the policy-induced state-action occupancy measures of the agent and the expert. Recently, Hazan et al. [21] or Zhang et al. [44] abstracted such problems as Concave (or general) Utility Reinforcement Learning (CURL), that aims at maximizing a concave function of the occupancy measure, over policies generating them. This calls for new algorithms since usual RL approaches do not readily apply.

Mean-field games (MFGs) have been introduced concurrently by Lasry and Lions [23] and Huang et al. [22]. They address sequential decision problems involving a large population of anonymous agents, with symmetric interests. They do so by considering the limiting case of a continuous distribution of agents, which allows reducing the original problem to the study of a single representative agent through its interaction with the full population. MFGs have been originally studied in Mathematics, with solutions typically involving coupled differential equations, but has recently seen a surge of interest in Machine Learning, e.g. [7, 19, 30, 42].

Our core contribution consists in showing that CURL is indeed a subclass of MFGs. We think this to be of interest to bridge together both communities and shed some light on connections between both fields. For example, we show an equivalence between the concavity of the CURL function and the monotonicity of the corresponding MFG (monotonicity being a sufficient condition to guarantee the existence of a Nash equilibrium). We also show an equivalence between the optimality conditions of the CURL problem (hence its maximizer) and the exploitability of the corresponding MFG (characterizing the Nash equilibrium). We also discuss algorithms for both fields. We show that the seminal approach of Hazan et al. [21] for CURL is nothing else than the classic Fictitious Play (FP) for MFGs [8, 30]. We also show that both are in fact instances of Frank-Wolfe [14] applied to the CURL problem. This readily provides the first (to the best of our knowledge) convergence rate for discrete-time FP applied to an MFG, consistent with the one hypothesized by Perrin et al. [30] while studying continuous-time FP. This connection between fields also allows applying MFGs algorithms to CURL. In a set of numerical illustrations, we apply the recently introduced Online Mirror Descent (OMD) for MFGs [29] to CURL, which appears to converge much faster than classic FP.

2 BACKGROUND AND RELATED WORKS

We write Δ_X the set of probability distributions over a finite set X and Y^X the set of applications from X to the set Y . For $f_1, f_2 \in \mathbb{R}^X$, we write the dot product $\langle f_1, f_2 \rangle = \sum_{x \in X} f_1(x) f_2(x)$. RL is usually formalized using Markov Decision Processes (MDPs). An MDP is a tuple $\{\mathcal{S}, \mathcal{A}, P, r, \gamma, \rho_0\}$ with \mathcal{S} the finite state space, \mathcal{A} the finite action space, $P \in \Delta_{\mathcal{S} \times \mathcal{A}}^{\mathcal{S}}$ the Markovian transition kernel, $r \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$ the (bounded) reward function, $\gamma \in (0, 1)$ the discount factor and $\rho_0 \in \Delta_{\mathcal{S}}$ the initial state distribution. A policy $\pi \in \Delta_{\mathcal{A}}^{\mathcal{S}}$ associates to each state a distribution over actions. Its (state-action) value function is $q_{\pi, r}(s, a) = \mathbb{E}_{\pi}[\sum_{t \geq 0} \gamma^t r(S_t, A_t) | S_0 = s, A_0 = a]$, with \mathbb{E}_{π} the expectation over trajectories induced by the policy π (and the dynamics P). RL aims at finding a policy maximizing the value for any state-action pair, or as considered here for initial states sampled according to ρ_0 . Writing $\mu_{0, \pi} = \rho_0 \pi \in \Delta_{\mathcal{S} \times \mathcal{A}}$ ($\mu_{0, \pi}(s, a) = \rho_0(s) \pi(a|s)$), the RL problem can thus be framed as (the factor $(1 - \gamma)$ accounting for later normalization) $\max_{\pi \in \Delta_{\mathcal{A}}^{\mathcal{S}}} J_{\Pi}(\pi)$ with

$$J_{\Pi}(\pi) = (1 - \gamma) \mathbb{E}_{(S, A) \sim \mu_{0, \pi}} [q_{\pi, r}(S, A)] = (1 - \gamma) \langle \mu_{0, \pi}, q_{\pi, r} \rangle.$$

Define the policy kernel P_{π} as $P_{\pi}(s', a' | s, a) = P(s' | s, a) \pi(a' | s')$. The value function is the unique fixed point of the Bellman equation [32], $q_{\pi, r} = r + \gamma P_{\pi} q_{\pi, r}$, or equivalently $q_{\pi, r} = (I - \gamma P_{\pi})^{-1} r$, with I the identity. We'll use the discounted occupancy measure induced by π , $\mu_{\pi} \in \Delta_{\mathcal{S} \times \mathcal{A}}$, defined as $\mu_{\pi}(s, a) = (1 - \gamma) \sum_{t \geq 0} \gamma^t \Pr(S_t = s, A_t = a | S_0 \sim \rho_0, \pi)$, or for short $\mu_{\pi} = (1 - \gamma)(I - \gamma P_{\pi}^{\top})^{-1} \mu_{0, \pi}$. With this additional notation, the RL objective function is

$$J_{\Pi}(\pi) = \langle \mu_{\pi}, r \rangle. \quad (1)$$

We clearly see the linearity of the RL objective in terms of the policy-induced occupancy measure (but not in terms of the policy). As it will be useful later, we also introduce the set $M \subset \Delta_{\mathcal{S} \times \mathcal{A}}$ of state-action distributions satisfying the Bellman flow (e.g., [44]):

$$M = \left\{ \mu \in \Delta_{\mathcal{S} \times \mathcal{A}} \mid \sum_{a \in \mathcal{A}} \mu(\cdot, a) = (1 - \gamma) \rho_0 + \gamma P^{\top} \mu \right\}. \quad (2)$$

For any $\mu \in \Delta_{\mathcal{S} \times \mathcal{A}}$, we write $\rho \in \Delta_{\mathcal{S}}$ its marginal over states:

$$\forall s \in \mathcal{S}, \rho(s) = \sum_{a \in \mathcal{A}} \mu(s, a). \quad (3)$$

For any policy $\pi \in \Delta_{\mathcal{A}}^{\mathcal{S}}$, we have $\mu_{\pi} \in M$, and for any $\mu \in M$, defining $\pi(a|s) = \frac{\mu(s, a)}{\rho(s)}$, we have $\mu_{\pi} = \mu$. Thus, the RL problem defined Eq. (1) is a linear program, $\max_{\pi \in \Delta_{\mathcal{A}}^{\mathcal{S}}} \langle \mu_{\pi}, r \rangle = \max_{\mu \in M} \langle \mu, r \rangle$.

2.1 Concave Utility Reinforcement Learning

Let $F \in \mathbb{R}^{\Delta_{\mathcal{S} \times \mathcal{A}}}$ be a concave function, CURL [21] aims at solving

$$\max_{\pi \in \Delta_{\mathcal{A}}^{\mathcal{S}}} F(\mu_{\pi}). \quad (4)$$

As discussed above, RL is a special case with $F(\mu) = \langle \mu, r \rangle$, but this representation also subsumes other relevant related problems. We present below a probably non-exhaustive list of them.

Exploration. Hazan et al. [21] formulated the pure exploration problem as finding the policy maximizing the entropy of the induced state occupancy measure, corresponding to $F(\mu) = -\langle \rho, \ln \rho \rangle$ (recall the definition of ρ from Eq. (3)). One could consider the entropy of the state-action occupancy measure, $F(\mu) = -\langle \mu, \ln \mu \rangle$, but it

would probably make less sense from an exploration perspective (we usually want to visit all states evenly, without caring about actions themselves). This could also be combined with the maximization of a cumulative reward, $F(\mu) = \langle \mu, r \rangle - \lambda \langle \rho, \ln \rho \rangle$, with λ balancing task resolution and exploration.

Divergence minimization. For a given f -divergence D_f , this task corresponds to $F(\mu) = -D_f(\mu || \mu_*)$ or $F(\mu) = -D_f(\rho || \rho_*)$, with μ_* (resp. ρ_*) a target occupancy measure. The first case encompasses many **imitation learning** approaches [16]. There, the target measure μ_* is given by trajectories of an expert policy to be imitated. The second case encompasses for example the recently introduced framework of **state-marginal matching** [24], that itself generalizes entropy maximization. The core difference with imitation learning is whether the occupancy measure of interest is directly accessible or only through samples. As noted by Zhang et al. [44], this approach can be generalized to any distance (e.g., Wasserstein), as long as it is convex in μ .

Constrained MDPs. Let $c \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$ be a cost function. A constrained MDP aims at maximizing the cumulative reward while keeping the cumulative cost below a threshold C : $\max_{\pi} \langle \mu_{\pi}, r \rangle$ such that $\langle \mu_{\pi}, c \rangle \leq C$. As observed by Zhang et al. [44], a relaxed formulation can be framed in the CURL framework, with $F(\mu) = \langle \mu, r \rangle - \lambda p(\langle \mu, c \rangle - C)$, where p is a convex penalty function.

Multi-objectif RL deals with the joint optimization of multiple rewards. Different approaches can be envisioned, some of them being a special case of CURL. Let $r^1 \dots r^K$ be K reward functions, and write $v_{\pi}^k = \langle \mu_{\pi}, r^k \rangle$. Consider $g : \mathbb{R}^K \rightarrow \mathbb{R}$ a concave function. Some multi-objective RL approaches [3, 10] aim at solving $\max_{\pi} g(v_{\pi}^1, \dots, v_{\pi}^K)$. As a composition of a linear and a concave functions, it is concave in μ_{π} , and thus a CURL problem.

Offline RL [25] aims at learning a policy from a fixed dataset, without interacting with the system. A classic approach consists in penalizing the learnt policy when deviating too much from the one that generated the data, through some implicit or explicit divergence between policies [25]. It would also make sense to penalize the learnt policy to reach states deviating from the ones in the dataset. If this approach was, up to our knowledge, never considered in the literature, it could also be framed as a CURL problem, with for example $F(\mu) = \langle \mu, r \rangle - \lambda D_f(\rho || \rho_{\text{data}})$. A similar and practical approach has been considered in the online RL setting [36].

2.2 Mean-Field Games

As explained in Sec. 1, MFGs are a continuous approximation of large population multi-agent problems. Thus, a core question is generally to know how good this approximation is, depending on the population size. Interestingly, we only need distributions of agents for drawing the connections to CURL, so we will not address this question further. Moreover, it means that, whenever we refer to a population or a distribution of agents, it indeed corresponds to the occupancy measure induced by a single policy (there is a single agent, *per se*). Generally speaking, both the dynamics and the reward can depend on the population, but for drawing connections to CURL we need only population-dependent rewards. The (restricted) setting we present next is inspired by e.g. [29, 30], other settings exist (some being discussed in more details in Sec. 4.3).

An MFG is here a tuple $\{\mathcal{S}, \mathcal{A}, P, \mathcal{R}, \gamma, \rho_0\}$. Everything is as in MDPs, except for the reward function $\mathcal{R} \in \mathbb{R}^{\mathcal{S} \times \mathcal{A} \times \Delta \mathcal{S}}$ that now also depends on the state-distribution of the population. For such a fixed distribution $\rho \in \mathbb{R}^{\mathcal{S}}$, we define the following MDP-like criterion:

$$\begin{aligned} J(\pi, \rho) &= (1 - \gamma) \mathbb{E}_\pi \left[\sum_{t \geq 0} \gamma^t \mathcal{R}(S_t, A_t, \rho) \mid S_0 \sim \rho_0 \right] \\ &= \langle \mu_\pi, \mathcal{R}(\cdot, \rho) \rangle. \end{aligned} \quad (5)$$

Exploitability. The exploitability $\phi(\pi)$ of a policy π quantifies the maximum gain a representative player can get by deviating its policy from the rest of the population still playing π ,

$$\phi(\pi) = \max_{\pi' \in \Delta_{\mathcal{S}^{\mathcal{A}}}} J(\pi', \rho_\pi) - J(\pi, \rho_\pi).$$

The maximizer of $J(\pi, \rho)$ for a fixed ρ is called a **best response** (optimal policy for the associated MDP).

Nash equilibrium. A Nash equilibrium is a policy satisfying $\phi(\pi) = 0$: there is nothing to gain from deviating from the population policy. The core problem of MFGs is to compute this Nash.

Monotonicity. Following Lasry and Lions [23], a game is said to be **monotone** if for any $\mu, \mu' \in \Delta_{\mathcal{S} \times \mathcal{A}}$, we have (with ρ and ρ' being defined through Eq. (3), again) $\langle \mu - \mu', \mathcal{R}(\cdot, \rho) - \mathcal{R}(\cdot, \rho') \rangle \leq 0$. This ensures existence of a Nash. If moreover the inequality is strict whenever $\mu \neq \mu'$, the game is said to be **strictly monotone**, which ensures uniqueness of the Nash [29].

Separable reward. This special case will be of interest later. The reward is said to be separable if it decomposes as $\mathcal{R}(s, a, \rho) = \bar{r}(s, a) + \tilde{r}(s, \rho)$. If in addition the following monotonicity condition holds, for all $\rho \neq \rho' \in \Delta_{\mathcal{S}}$, $\langle \rho - \rho', \tilde{r}(\cdot, \rho) - \tilde{r}(\cdot, \rho') \rangle \leq 0$ (resp. < 0), then the game is monotone (resp. strictly monotone) [29].

We will also consider less usual and slightly more general MFGs (sometimes referred to as Extended MFGs in the literature [18]), where the reward now depends on the state-action distribution, rather than on the state-distribution, $\mathcal{R} \in \mathbb{R}^{\mathcal{S} \times \mathcal{A} \times \Delta_{\mathcal{S} \times \mathcal{A}}}$. The MDP-like criterion is thus $J(\pi, \mu) = \langle \mu_\pi, \mathcal{R}(\cdot, \mu) \rangle$. Similarly to the previous case, we will say that the game is monotone (resp. strictly) if for any $\mu \neq \mu' \in \Delta_{\mathcal{S} \times \mathcal{A}}$, we have $\langle \mu - \mu', \mathcal{R}(\cdot, \mu) - \mathcal{R}(\cdot, \mu') \rangle \leq 0$ (resp. < 0). Existence or uniqueness of the Nash through this monotonicity is not readily available from the literature, as far as we know, at least in an MDP setting. Yet, it will be a direct consequence of Sec. 3.

2.3 Related Works

As far as we know, this work is the first to connect CURL with MFGs. We briefly review relevant works of each of these fields.

CURL. All problems subsumed by CURL benefit from a large body of literature. For example, many articles address the problem of exploration in RL, not necessarily through the lens of entropy maximization (even though links can be drawn between predicting-error-based exploration and entropy maximization [24]). As stated before, many IL algorithms minimize a divergence between occupancy measures. Yet, they do not do it directly, and rather optimize a variational bound (to get a saddle-point problem involving a difference of expectation, to avoid estimating the target density). The approach of Lee et al. [24] for marginal matching minimizes a KL divergence differently, but the resulting algorithm is equivalent to the one of Hazan et al. [21], which tackles more general problems. There are two approaches that address the general CURL problem:

the one of Hazan et al. [21] and the one of Zhang et al. [44], discussed more extensively in Sec. 4. Concurrently to our work, Zahavy et al. [43] also tackle the general CURL problem, by transforming it into a saddle-point problem, thanks to the Legendre-Fenchel transform of the concave objective. They then adapt the meta-algorithm of Abernethy and Wang [2] and show that many algorithms handling special cases of CURL can be derived from it, yet relying on an explicit form of the convex conjugate of F . Our approach is different and complementary: we frame CURL as an MFG.

MFGs. There is also a large body of literature on MFGs. We will focus on the part more closely related to machine learning. A seminal paper combining MFG with learning was proposed by Yang et al. [42] who analyze the convergence of mean field Q-learning and Actor-Critic algorithms to Nash. Guo et al. [19] also combine Q-learning with MFGs, through a fixed-point approach. Elie et al. [13] studies the propagation of errors due to learning and its effect on the Nash equilibrium. Perrin et al. [30] studies FP, both theoretically and empirically, when combined with learning. Perolat et al. [29] introduces a mirror descent approach to MFG. These are only a few examples among others [26, 35, 41]. Connections between MFGs and GANs [7] or Deep Learning architectures [12] have also been established. We think that an interesting aspect of our contribution is that many progress made on the MFGs side could be readily beneficial to CURL, thanks to what will be presented in Sec. 4.3 (but up to a compatible setting, this will be discussed in Sec. 4). The converse may also be true: advances in CURL can benefit (possibly a subclass of) MFGs. Another relevant part of the MFG literature are potential MFGs (e.g., [8, 23]), where the reward function derives from a potential function. Results similar to those that we state in Sec. 3 appear there (link between monotonicity and concavity, maximizer and Nash), but within a different paradigm (notably, not within the MDP setting, possibly deterministic dynamics, specified through differential equations, finite horizon or even static case, etc.), and with no link to (CU)RL (e.g., [9, Prop. 1.10, Lemma 5.72]).

3 CURL IS AN MFG

Now, we frame the CURL problem (4) as an MFG. We handle first the general case $F(\mu)$, and then the special case $F(\mu) = \langle \mu, r \rangle + G(\rho)$, with possibly $r = 0$, as it has a specific structure representative of many problems depicted in Sec. 2.1.

General Case. Here, we consider the general CURL problem (4). We define a potential MFG with the same dynamics, discount and initial distribution, but with reward deriving from the potential F :

$$\forall \mu \in \Delta_{\mathcal{S} \times \mathcal{A}}, \mathcal{R}(\cdot, \mu) = \nabla F(\mu) \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}. \quad (6)$$

The next results link the concavity of F to the monotonicity of the game, and maximizers of CURL to Nash equilibria of the MFG.

PROPERTY 1. *The MFG defined Eq. (6) is (strictly) monotone if and only if F is (strictly) concave.*

PROOF. From basic convex optimization principles, we have

$$\begin{aligned} F \text{ concave} &\stackrel{(a)}{\Leftrightarrow} \forall \mu_1, \mu_2 \in \Delta_{\mathcal{S} \times \mathcal{A}}, \langle \mu_2 - \mu_1, \nabla F(\mu_2) - \nabla F(\mu_1) \rangle \leq 0 \\ &\stackrel{(b)}{\Leftrightarrow} \forall \mu_1, \mu_2 \in \Delta_{\mathcal{S} \times \mathcal{A}}, \langle \mu_2 - \mu_1, \mathcal{R}(\cdot, \mu_2) - \mathcal{R}(\cdot, \mu_1) \rangle \leq 0 \\ &\stackrel{(c)}{\Leftrightarrow} \text{the MFG is monotone,} \end{aligned}$$

with (a) by [27, Thm. 2.1.3], (b) by Eq. (6) and (c) by definition. The equivalence between strict concavity and strict monotonicity is obtained by replacing inequalities by strict ones for $\mu_1 \neq \mu_2$. \square

THEOREM 1. *Assume that F is concave. Any maximizer of the CURL problem (4) is a Nash equilibrium of the MFG defined through reward (6), and conversely any Nash is a maximizer.*

PROOF. We have that $\max_{\pi \in \Delta_{\mathcal{A}}^S} F(\mu_\pi) = \max_{\mu \in M} F(\mu)$, with M defined in Eq. (2). The later is a concave program with linear constraints, so concavity of F ensures the existence of a maximizer, satisfying the optimality condition. Also, as seen in Sec. 2, for any policy π , $\mu_\pi \in M$, and for any $\mu \in M$, there is an associated policy (the conditional on actions, $\mu(s, a)/\rho(s)$). Therefore, we have

$$\begin{aligned} \pi \text{ maximizer} &\stackrel{(a)}{\Leftrightarrow} \langle \nabla F(\mu_\pi), \mu' - \mu_\pi \rangle \leq 0, \forall \mu' \in M \\ &\stackrel{(b)}{\Leftrightarrow} \langle \mathcal{R}(\cdot, \mu_\pi), \mu_{\pi'} - \mu_\pi \rangle \leq 0, \forall \pi' \in \Delta_{\mathcal{A}}^S \\ &\Leftrightarrow \langle \mu'_{\pi'}, \mathcal{R}(\cdot, \mu_\pi) \rangle \leq \langle \mu_\pi, \mathcal{R}(\cdot, \mu_\pi) \rangle, \forall \pi' \in \Delta_{\mathcal{A}}^S \\ &\stackrel{(c)}{\Leftrightarrow} J(\pi', \mu_\pi) \leq J(\pi, \mu_\pi), \forall \pi' \in \Delta_{\mathcal{A}}^S \\ &\stackrel{(d)}{\Leftrightarrow} \phi(\pi) = 0 \Leftrightarrow \pi \text{ is Nash,} \end{aligned}$$

with (a) by optimality conditions, (b) by using the associated policy and Eq. (6), (c) by def. of J and (d) by def. of exploitability. \square

Interestingly, the proof of Thm. 1 shows an equivalence between the optimality conditions for CURL (that provides a global optimum, even though $F(\mu_\pi)$ is not concave in π , thanks to the relationship between $\Delta_{\mathcal{A}}^S$ and M) and a null exploitability for MFG. As stated Sec. 2.2, existence or uniqueness of the Nash for this class of MFGs from (strict) monotonicity is not readily available in the literature, in this MDP setting. Yet, by Prop. 1 monotonicity and concavity are equivalent, by Thm. 1 the set of CURL maximizers and that of MFG Nash equilibria are equal, so the result holds readily (uniqueness in M thanks to strict concavity implying uniqueness in $\Delta_{\mathcal{A}}^S$ here).

A Relevant Special Case. We consider CURL (4) when F takes the form (ρ as in Eq. (3))

$$F(\mu) = \langle \mu, r \rangle + G(\rho), \quad (7)$$

with r being possibly null and G being a differentiable potential function. This is representative of a number of examples of Sec. 2.1. For this, we define the MFG with the following separable reward,

$$\mathcal{R}(s, a, \rho) = \bar{r}(s, a) + \tilde{r}(s, \rho) \text{ with } \bar{r} = r \text{ and } \tilde{r}(\cdot, \rho) = \nabla G(\rho). \quad (8)$$

We have the counterparts of Prop. 1 and Thm. 1.

PROPERTY 2. *The MFG defined Eq. (8) is (strictly) monotone if and only if G is (strictly) concave.*

PROOF. This is similar to the proof of Prop. 1. We have

$$\begin{aligned} G \text{ concave} &\Leftrightarrow \forall \rho_1, \rho_2 \in \Delta_S, \langle \rho_1 - \rho_2, \nabla G(\rho_1) - \nabla G(\rho_2) \rangle \leq 0 \\ &\Leftrightarrow \forall \rho_1, \rho_2 \in \Delta_S, \langle \rho_1 - \rho_2, \tilde{r}(\cdot, \rho_1) - \tilde{r}(\cdot, \rho_2) \rangle \leq 0. \end{aligned}$$

The equivalence between the last condition and (strict) monotonicity follows from [29, Lemma 2]. \square

THEOREM 2. *Assume that G is concave. Any maximizer of the CURL problem (4) with F defined in Eq. (7) is a Nash equilibrium of the MFG defined through reward (8), and conversely.*

PROOF. The proof follows globally the same lines as the one of Thm. 1. As G is concave in ρ , and ρ being the marginal of μ , F is concave in μ . Therefore, satisfying the optimality conditions provides a global optimizer. Before proceeding further, notice that $\nabla F(\mu) = r + \nabla_\mu G(\rho)$. Using the chain rule, we have

$$\frac{\partial G(\rho)}{\partial \mu(s, a)} = \frac{\partial G(\rho)}{\partial \rho(s)} \frac{\partial \rho(s)}{\partial \mu(s, a)} = \frac{\partial G(\rho)}{\partial \rho(s)} \frac{\partial \sum_{a' \in \mathcal{A}} \mu(s, a')}{\partial \mu(s, a)} = \frac{\partial G(\rho)}{\partial \rho(s)}.$$

Therefore, we have

$$\begin{aligned} \langle \mu', \nabla_\mu G(\rho) \rangle &= \sum_{s, a} \mu'(s, a) \frac{\partial G(\rho)}{\partial \rho(s)} = \sum_s \frac{\partial G(\rho)}{\partial \rho(s)} \sum_a \mu'(s, a) \\ &= \sum_s \frac{\partial G(\rho)}{\partial \rho(s)} \rho'(s) = \langle \rho', \nabla G(\rho) \rangle. \end{aligned}$$

Hence, we have π maximizer

$$\begin{aligned} &\Leftrightarrow \langle \nabla F(\mu_\pi), \mu' - \mu_\pi \rangle \leq 0 && \forall \mu' \in M \\ &\Leftrightarrow \langle r + \nabla_\mu G(\rho_\pi), \mu' - \mu_\pi \rangle \leq 0 && \forall \mu' \in M \\ &\Leftrightarrow \langle \mu', r \rangle + \langle \rho', \nabla G(\rho_\pi) \rangle \leq \langle \mu_\pi, r \rangle + \langle \rho_\pi, \nabla G(\rho_\pi) \rangle && \forall \mu' \in M \\ &\Leftrightarrow J(\pi', \mu_\pi) \leq J(\pi, \mu_\pi) && \forall \pi' \in \Delta_{\mathcal{A}}^S \\ &\Leftrightarrow \phi(\pi) = 0 \Leftrightarrow \pi \text{ is Nash.} \end{aligned}$$

This proves the stated result. \square

One should be meticulous about uniqueness properties in this case. We have shown in Prop. 2 the equivalence between the strict concavity of G and the strict monotonicity of the game, hence uniqueness of the Nash. Yet, what is unique in this case is the population ρ_π , but not necessarily the joint distribution μ_π or the policy π (e.g., an MDP with two actions having the same effect).

This equivalence between CURL and MFG is interesting, because any algorithm designed for this setting of MFGs can be readily applied to CURL, and conversely, at least when the MFG is potential and when the dynamics does not depend on the population.

4 ALGORITHMS

Before presenting algorithms, we discuss how to measure their progresses. From a CURL viewpoint, it is natural to measure how F increases with iterates. Yet, the maximum is usually not known a priori. From an MFG perspective, it may be more natural to use exploitability, that we know should be zero at optimality. Yet, it is more costly, as its evaluation requires computing a best response. Now, we will show that the seminal approach to CURL is indeed FP for MFGs, and connect it to Frank-Wolfe. Then, we will discuss briefly other algorithms specific to either CURL or MFGs.

4.1 On Hazan et al. [21] Approach, Fictitious Play and Frank-Wolfe

Hazan et al. [21] approach. As far as we know, the first approach that addresses the general CURL problem was proposed by Hazan et al. [21]. They consider the case $F(\mu) = G(\rho)$ (so a special case of what was studied in Sec. 3, but the approach could readily be extended to a general $F(\mu)$). For a set of k policies $C = (\pi_0, \dots, \pi_{k-1})$

Algorithm 1: Hazan et al. [21] algorithm.

Input: Step size η , number of iterations T , initial policy π_0
Set $C_0 = (\pi_0)$, $\alpha_0 = 1$, $\pi_{\text{mix},0} = (\alpha_0, C_0)$;
for $t = 0, \dots, T - 1$ **do**
 Compute π_{t+1} , optimal for reward $\nabla G(\rho_{\pi_{\text{mix},t}}) \in \mathbb{R}^{\mathcal{S}}$;
 Update $\pi_{\text{mix},t+1} = (\alpha_{t+1}, C_{t+1})$ with $C_{t+1} = (C_t, \pi_{t+1})$
 and $\alpha_{t+1} = ((1 - \eta)\alpha_t, \eta)$;
return $\pi_{\text{mix},T}$

and weights $\alpha \in \Delta_{\{0, \dots, k-1\}}$, they define the non-stationary policy π_{mix} as the one that samples a policy π_i with probability α_i in the initial state, and this policy is used for the whole trajectory. The induced occupancy measure satisfies $\rho_{\pi_{\text{mix}}} = \sum_{i=0}^{k-1} \alpha_i \rho_{\pi_i}$. Their approach is inspired by Frank-Wolfe [14], and is depicted in Alg. 1 (here assuming a known model). The mixing policy is initialized with some given policy. At iteration t , a new policy being the best response to the MDP defined through the reward $\nabla G(\rho_{\pi_{\text{mix},t}})$ is added to the mixture, and the set of weights is updated according to a predefined learning rate. Therefore, the CURL problem is reduced to a sequence of MDP problems. Notice that what characterizes the optimality of the policy is the state occupancy measure it induces. Therefore, instead of returning the non-stationary policy $\pi_{\text{mix},T}$, one could return the stationary policy $\pi_{\text{sta},T}(a|s) = \sum_{t=0}^T \alpha_T(t) \rho_{\pi_t}(s) \pi_t(a|s) (\sum_{t=0}^T \alpha_T(t) \rho_{\pi_t}(s))^{-1}$, as it induces the same state occupancy measure. This implies that Alg. 1 is indeed Fictitious Play, as explained further later. Notice also that on an abstract way, the approach of Lee et al. [24] is exactly the same: it differs from [21] in the kind of function G considered (specifically a KL divergence for [24]), and on how approximations are done without knowledge of the domain (something we will not address here, but that comes with guarantees for Hazan et al. [21]).

Algorithm 2: Fictitious Play [30].

Input: number of iterations T , initial policy π_0
Set $\bar{\mu}_0 = \mu_{\pi_0}$;
for $t = 0, \dots, T - 1$ **do**
 Compute $\pi_{t+1} \in \operatorname{argmax}_{\pi \in \Delta_{\mathcal{A}}^{\mathcal{S}}} J(\pi, \bar{\mu}_t)$;
 Update $\bar{\mu}_{t+1} = \frac{1}{t+2} \mu_{\pi_{t+1}} + \frac{t+1}{t+2} \bar{\mu}_t$;
return $\bar{\pi}_T : \bar{\pi}_T(a|s) = \frac{\bar{\mu}_T(s,a)}{\bar{\rho}_T(s)}$

Fictitious Play. It is originally an iterative algorithm for repeated games [33], where each agent plays optimally according to an empirical expectation of past observed strategies. It has been adapted to MFGs [8, 30] with the approach summarized in Alg. 2. This is for the case where $\mathcal{R} \in \mathbb{R}^{\mathcal{S} \times \mathcal{A} \times \Delta_{\mathcal{S} \times \mathcal{A}}}$, but one just needs to replace $J(\pi, \bar{\mu}_t)$ by $J(\pi, \bar{\rho}_t)$ (with $\bar{\rho}$ still the marginal of $\bar{\mu}$, as defined Eq. (3)) to cover the other case. The algorithm is initialized for some π_0 , and at each iteration a best response is computed by maximizing $J(\pi, \bar{\mu}_t)$ (Eq. (5)), with $\bar{\mu}_t$ the average of all past state-action occupancy measures (associated to all past best responses). The output policy is the stationary one having $\bar{\mu}_T$ as associated occupancy measure (see Sec. 2). With this, we can easily see that Alg. 1 is indeed an FP approach, equivalent to Alg. 2. With the

choice $\eta = \frac{1}{t+2}$, we have that $\rho_{\pi_{\text{mix},t}} = \bar{\rho}_t$. With the MFG defined in Sec. 3, that is, with reward $\mathcal{R}(\cdot, \rho) = \nabla G(\rho)$ (Eq. (8)), maximizing $J(\pi, \bar{\rho}_t)$ is equivalent to solving the MDP with reward $\nabla G(\rho_{\text{mix},t})$. Eventually, both policies $\pi_{\text{mix},T}$ (non-stationary) and $\bar{\pi}_T$ (stationary) induce the same state occupancy measure, and thus the same solution to the problem.

Frank-Wolfe. Now, we take advantage of the relationship between CURL and MFGs to provide a new interpretation of FP for (potential) MFGs as equivalent to Frank-Wolfe (FW) applied to the original optimization problem. This notably allows to provide the first convergence rate for discrete-time FP in this setting. As seen before, CURL (4) can be defined solely in terms of occupancy measure, $\max_{\mu \in M} F(\mu)$. The set M (2) being defined as a set of linear constraints, it is a convex set. With F concave and M convex, we can readily apply FW. With $\mu_0 \in M$, for all $t \geq 0$,

$$\mu_{t+1/2} \in \operatorname{argmax}_{\mu \in M} \langle \mu, \nabla F(\mu_t) \rangle \text{ and } \mu_{t+1} = (1 - \eta_{t+1})\mu_t + \eta_{t+1}\mu_{t+1/2}.$$

First, one searches for the element of M which is the most collinear with the gradient. Then, this is used to update the estimate, so as to stay in the domain, $\mu_{t+1} = (1 - \eta_{t+1})\mu_t + \eta_{t+1}\mu_{t+1/2}$. With $\eta_t = \eta$, we retrieve the occupancy measures induced by the mixing policy of [21], and with $\eta_t = \frac{1}{t+1}$ we obtain the averaging of FP. Now, for the inner optimization problem, observe that

$$\max_{\mu \in M} \langle \mu, \nabla F(\mu_t) \rangle = \max_{\pi \in \Delta_{\mathcal{A}}^{\mathcal{S}}} \langle \mu_{\pi}, \nabla F(\mu_t) \rangle,$$

which corresponds exactly to the best response computed in both Algs. 1 and 2. So, both are indeed special instances of FW. This is not new for the approach of Hazan et al. [21], who refer explicitly to FW (but do not make any connection to FP or MFGs, or games in general), but this is new regarding FP for (potential) MFGs, as far as we know. This leads to the following result.

THEOREM 3. Consider a monotone potential MFG, that is such that $\mathcal{R}(\cdot, \mu) = \nabla F(\mu)$ with F concave (Prop. 1). Write $R = \max_{\mu, \mu' \in M} \|\mu - \mu'\|$ the diameter of M . Let $\mu_* \in M$ be a maximizer of F (or equivalently by Thm. 1, let $\pi_*(a|s) = \mu_*(s, a) / \rho_*(s)$ be a Nash of the MFG). Assume that the reward is β -Lipschitz in the sense that $\forall \mu, \mu' \in \Delta_{\mathcal{S} \times \mathcal{A}}$, $\|\mathcal{R}(\cdot, \mu) - \mathcal{R}(\cdot, \mu')\| \leq \beta \|\mu - \mu'\|$. Generally speaking, the exploitability is bounded by the suboptimality. For any $\pi \in \Delta_{\mathcal{A}}^{\mathcal{S}}$:

$$\phi(\pi) \leq F(\mu_*) - F(\mu_{\pi}) + R\sqrt{2\beta(F(\mu_*) - F(\mu_{\pi}))}.$$

Specifically, by running FW/FP for T iterations, with $\eta_t = \frac{2}{t+1}$,

$$\phi(\pi_T) \leq \frac{2\beta R^2}{\sqrt{T+1}} + \frac{2\beta R^2}{T+1} \text{ with } \pi_T(a|s) = \frac{\mu_T(s, a)}{\rho_T(s)}.$$

PROOF. First, notice that the β -Lipschitzness of the reward is equivalent to the classic β -smoothness of the function F . For the first result, we will exploit the equivalence between upper-bounding the exploitability $\phi(\pi)$ and upper-bounding the optimality condition $\langle \nabla F(\mu_{\pi}), \mu' - \mu_{\pi} \rangle$ (for arbitrary $\mu' \in M$). Indeed, from the proof of Thm. 1, let $\epsilon > 0$ and $\pi \in \Delta_{\mathcal{A}}^{\mathcal{S}}$ be an arbitrary policy,

$$\forall \mu' \in M, \langle \nabla F(\mu_{\pi}), \mu' - \mu_{\pi} \rangle \leq \epsilon \Leftrightarrow \phi(\pi) \leq \epsilon. \quad (9)$$

For any $\mu' \in M$, with a simple decomposition, we have that

$$\langle \nabla F(\mu_{\pi}), \mu' - \mu_{\pi} \rangle = \langle \nabla F(\mu_*), \mu' - \mu_{\pi} \rangle + \langle \nabla F(\mu_{\pi}) - \nabla F(\mu_*), \mu' - \mu_{\pi} \rangle.$$

We now upper-bound each of these two terms. For the first one:

$$\begin{aligned} \langle \nabla F(\mu_*), \mu' - \mu_\pi \rangle &= \langle \nabla F(\mu_*), \mu' - \mu_* \rangle + \langle \nabla F(\mu_*), \mu_* - \mu_\pi \rangle \\ &\stackrel{(a)}{\leq} \langle \nabla F(\mu_*), \mu_* - \mu_\pi \rangle \stackrel{(b)}{\leq} F(\mu_*) - F(\mu_\pi), \end{aligned}$$

with (a) by optimality conditions and (b) by concavity of F . For the second term, we have the following, using notably a classic result from convex optimization for β -smooth functions,

$$\begin{aligned} &\langle \nabla F(\mu_\pi) - \nabla F(\mu_*), \mu' - \mu_\pi \rangle \\ &\stackrel{(a)}{\leq} \|\nabla F(\mu_\pi) - \nabla F(\mu_*)\| \|\mu' - \mu_\pi\| \stackrel{(b)}{\leq} R \|\nabla F(\mu_\pi) - \nabla F(\mu_*)\| \\ &\stackrel{(c)}{\leq} R \sqrt{2\beta (\langle \nabla F(\mu_*), \mu_\pi - \mu_* \rangle + F(\mu_*) - F(\mu_\pi))} \\ &\stackrel{(d)}{\leq} R \sqrt{2\beta (F(\mu_*) - F(\mu_\pi))}, \end{aligned}$$

with (a) by Cauchy-Schwartz, (b) by def. of the diameter, (c) by [6, Lemma 3.5] applied to $-F$ and (d) by optimality conditions. Combining both bounds with the equivalence result of Eq. (9), we get the first stated result. For the second inequality, the set M being convex, and F being concave, we can apply the classic convergence result for FW (e.g., [6, Thm. 3.8]) to obtain

$$F(\mu_*) - F(\mu_T) \leq \frac{2\beta R^2}{T+1}.$$

Plugging this into the exploitability bound allows concluding. \square

Perrin et al. [30] showed that continuous-time FP (time of iterations, so a theoretical construct) enjoys a rate $\phi(T) = O(T^{-1})$ for monotone MFGs. They conjectured a rate $\phi(T) = O(T^{-1/2})$ for discrete-time FP. This is exactly what provides Thm. 3, yet under stronger assumptions (the MFG needs to be potential, and the reward β -Lipschitz). The reward smoothness may not be always satisfied, as for the entropy example. Yet, a smoothed objective can also be considered (e.g., a smoothed version $\mathcal{H}_\epsilon(\rho) = -\langle \rho, \ln(\rho + \epsilon) \rangle$ of the entropy [21]). Notice also that we prove that the exploitability is upper-bounded by the suboptimality. Thus, any better bound on FW (e.g., based on stronger assumptions) would translate readily into a faster decrease of the exploitability. FW has been extended to saddle-point problems, and in the bilinear case shown to be equivalent to FP [17, 20]. However, our connection between FW and FP for MFGs is new, to the best of our knowledge.

4.2 Online Mirror Descent

Algorithm 3: Online Mirror Descent [29].

Input: number of iterations T , initial policy π_0 , parameter α
Set $Q_0 = 0$;
for $t = 0, \dots, T - 1$ **do**
 Compute $q_{\pi_t, \mathcal{R}(\cdot, \mu_{\pi_t})}$;
 Update $Q_{t+1} = Q_t + \alpha q_{\pi_t, \mathcal{R}(\cdot, \mu_{\pi_t})}$;
 Define $\pi_{t+1} = \text{softmax}(Q_{t+1})$;
return π_T

To emphasize the potential benefits of addressing the CURL problem under the MFG point of view, we specifically consider the Online Mirror Descent algorithm (Alg. 3), recently adapted to

solve MFGs [29]. It is empirically faster than FP on various MFGs and comes with additional advantages. The algorithm is initialized with a policy π_0 . At each iteration, the policy π_t being given, one can compute or estimate the associated occupancy measure μ_{π_t} (as before, considering a state occupancy measure just amounts to replacing μ by ρ). This being fixed, one can compute the state-action value function $q_{\pi_t, \mathcal{R}(\cdot, \mu_{\pi_t})} \in \mathbb{R}^{S \times \mathcal{A}}$ of the policy π_t with the reward being $r = \mathcal{R}(\cdot, \mu_{\pi_t})$. This is a first advantage compared to FP: one simply needs to perform a policy evaluation step instead of computing a best response (that is, solving a full MDP). Then, the scaled value function is accumulated in $Q_{t+1} \in \mathbb{R}^{S \times \mathcal{A}}$. This is a second advantage compared to FP: it might be easier to accumulate value functions rather than occupancy measures, especially in an approximate case (not covered here, again). Eventually, the policy is simply the softmax of the scaled sum of past value functions (hence the parameter α can be understood as an inverse temperature). Perolat et al. [29] show the convergence of a continuous-time variation of this algorithm (here, the time is the iteration time t , so this corresponds to a theoretical abstraction of Alg. 3). Notice that this OMD for MFGs can be seen as a generalization of Mirror-Descent Policy Iteration [15] (retrieved for linear F), which is an abstraction of efficient deep RL algorithms such as TRPO [34], MPO [1] or M-RL [38], and which comes with strong theoretical guarantees in the RL setting [37]. Thus, provided a good density estimator (e.g., a normalizing flow [28]), this could be a good basis for a deep MFG agent, but this is beyond the scope of the current paper.

4.3 Other Algorithms

CURL-related algorithms. Apart from [21] discussed above, Zhang et al. [44] also address the general problem, by introducing a policy gradient approach to CURL. It is a gradient ascent of $F(\mu_\pi)$ according to π . Even though F is not concave in π , they provide a global convergence result for their approach, exploiting a so-called “hidden concavity” (roughly speaking, exploiting the fact that $\Delta_{\mathcal{A}}^S$ and M are in bijection, which is true at least for the tabular case). The naive gradient $\nabla_\pi F(\mu_\pi)$ (obtained using the chain-rule) cannot be estimated in general. They thus exploit the Legendre-Fenchel transform of $-F$ to express this gradient as the solution of a saddle-point problem, involving quantities that are easier to estimate. As such, their approach can be considered as a reduction from CURL to a sequence of (zero-sum) games. We will focus for the rest of this paper on the simpler case where CURL is reduced to a sequence of MDP problems, at most. The concurrent work of Zahavy et al. [43] provides a meta-algorithm for solving the CURL problem, that can be instantiated by combining regret minimizing algorithms (one for the best response, one for computing the reward as a function of all past occupancy measures). They also retrieve FW as a special case. Contrary to OMD, they require to compute best responses (except if relying on non-stationary, finite horizon RL), and the reward to be considered at each iteration depends on all past occupancy measures (instead of the last one only). It also relies on the conjugate of F .

MFG-related algorithms. We covered the MFG literature in Sec. 2, especially the part dealing with MDPs and learning. However, many of these approaches, e.g. [4, 5, 19, 35, 40], rely on a fixed-point setting, and make an implicit ergodic assumption, effectively getting ride of the initial state distribution. We explain briefly why, ignoring

the additional dependency of the dynamics to the population usually considered. For $\rho \in \Delta_{\mathcal{S}}$, define $\Gamma_1(\rho) \in \Delta_{\mathcal{A}}$ the optimal policy for the reward $\mathcal{R}(\cdot, \rho)$. Slightly abusing notations, write P_π the policy-induced state-transition kernel ($P_\pi(s'|s) = \sum_a \pi(a|s)P(s'|s, a)$). For $\pi \in \Delta_{\mathcal{A}}$ and $\rho \in \Delta_{\mathcal{S}}$, define $\Gamma_2(\pi, s) = P_\pi^\top \rho \in \Delta_{\mathcal{S}}$, the state distribution obtained by starting from ρ and applying π for one step. Then, define $\Gamma(\rho) = \Gamma_2(\Gamma_1(\rho), \rho) \in \Delta_{\mathcal{S}}$. These works make structural assumptions about the MDP so that Γ is a contraction (which do not hold for finite MDPs [11]), and the Nash equilibrium is then the fixed point of this operator, $\rho_* = \Gamma(\rho_*)$. First, notice that ρ_* is a stationary distribution, not a discounted occupancy measure, so $\rho_* \notin M$. So, this does not provide a solution to the defined CURL problem. Second, this notion of Nash equilibrium is different from the one we considered. We do require a representative player to not be able to gain something when both this player and the population start from the initial state distribution. With the fixed-point approach, it is required from the representative player to not be able to gain something when the population is already at equilibrium (that is, has reached its asymptotic behavior). For these reasons, applying the family of related algorithms to the CURL problem is not obvious¹. Yet, algorithms framed for the setting considered here (and not introduced in this paper) can already be of interest for CURL, as illustrated empirically next.

5 NUMERICAL ILLUSTRATION

This section aims at providing additional evidences of the benefits from connecting CURL and MFGs, through a set of numerical illustrations. We consider an environment simple enough to be visualized and instantiate on it a number of special cases of CURL, depicted in Sec. 2.1. We consider the model to be known and compute best responses exactly (or do exact policy evaluation). Our aim is to compare a representative algorithm for CURL, the one from Hazan et al. [21] that we have shown to be equivalent to FP in an MFG context, to an algorithm only considered in the MFG setting, namely OMD [29]. What we do not address here is the approximate case. Yet, each of the problems considered is not straightforward, even with a known model. Moreover, we think that density estimation methods considered in the literature [21, 24] for Hazan et al. [21] approach could be applied to OMD (see also discussion Sec. 4.2).

Fig. 1 shows the considered environment and setting. The environment is a simple four-rooms problem (four actions, up, down, left and right; moving towards a wall does not change the state; deterministic dynamics), the initial distribution being a Dirac on the upper-left corner. We consider four settings: entropy maximization (the log-density of the population corresponding to a uniform policy being given in Fig. 1 as a baseline), marginal matching (using a KL-divergence, the target distribution being in Fig. 1; notice that it is not achievable by an agent, the support being not connected), a constrained MDP (cost and reward are depicted in Fig. 1), and multi-objective RL (with three rewards depicted Fig. 1). For each problem, we compare FP (that is, the seminal approach of [21] to CURL) to OMD (specific to MFGs so far), both regarding the optimized function (CURL viewpoint) and the exploitability (MFG viewpoint), as a function of the number of policy evaluations (FP is run for 300

¹Alternatively, we could define a variant of the CURL problem with stationary rather than discounted occupancy measures, and this family of algorithms could apply.




Figure 1: Reading order: (a) the considered environment initial state in yellow, walls in white); (b) the log-density of a uniform policy (to illustrate entropy maximization); (c) the target log-density for marginal matching; (d) the constrained MDP (reward in yellow, constraints in light blue); (e) the three rewards of multi-objective RL.

iterations, and then OMD is run to have approximately the same number of policy evaluations). We also show the population (that is, occupancy measure) computed by each approach.

For FP, we used the FP rate $\eta_t = \frac{1}{t+1}$. We also tried FW rate ($\eta_t = \frac{2}{t+1}$), which provides similar results, and a constant rate [21] ($\eta_t = \eta$), which depends on the value of η and is never better than FP or FW. This is to be expected, this constant rate was designed to offer guarantees in the approximate case, while we consider the exact one. For OMD, in all experiments, we use $\alpha = 0.01$.

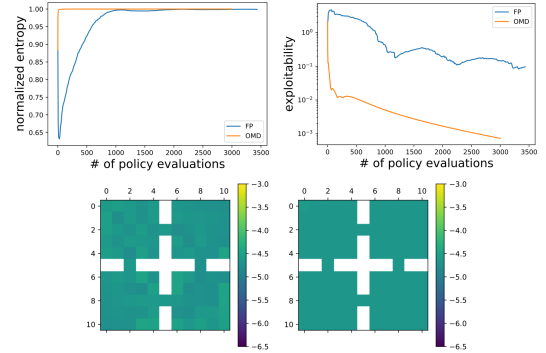


Figure 2: Entropy maximization. Reading order: (a) normalized entropy; (b) exploitability; (c) log-density for the FP population; (d) log-density for the OMD population.

Entropy maximization. The optimized function is $F_{\text{ent}}(\mu) = G_{\text{ent}}(\rho) = -\langle \rho, \ln \rho \rangle$, with related reward $\mathcal{R}_{\text{ent}}(\cdot, \mu) = \nabla G_{\text{ent}}(\rho) = -\ln \rho - 1$. Results are presented Fig. 2 (the entropy is normalized by $\ln(|\mathcal{S}|)$, the entropy of a uniform distribution). We can observe that OMD is much faster than FP both at increasing the entropy and decreasing the exploitability. At the final iteration, the entropies are hardly distinguishable numerically, but one can observe on the log-density plots that the population of OMD is more uniform (the scale is the same as for the uniform policy in Fig. 1.b, to ease comparison), which is consistent with the much lower exploitability.

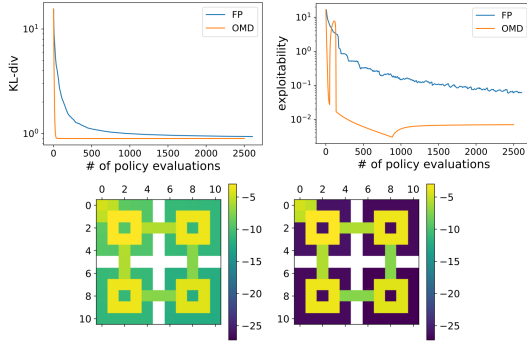


Figure 3: Marginal matching. Reading order: (a) divergence; (b) exploitability; (c) log-density for the FP population; (d) log-density for the OMD population.

Marginal matching. Here, the optimized function is $F_{\text{mm}}(\mu) = G_{\text{mm}}(\rho) = -\text{KL}(\rho||\rho_*)$, with ρ_* depicted in Fig. 1, and the corresponding reward is $\mathcal{R}_{\text{mm}}(\cdot, \mu) = \nabla G_{\text{mm}}(\rho) = \ln \rho_* - \ln \rho - 1$. Results are presented Fig. 3. Purposely, the target density is chosen such that it is not reachable by an agent. This illustrates the interest of the exploitability, which is zero at optimality, while we don't know a priori what the best solution is in terms of divergence. Here again, OMD is much faster at both optimizing the objective function and decreasing the exploitability (it is not monotonic for OMD, but nothing says it should be). We also observe from the log-density of the populations that OMD finds arguably a better solution, given the allowed number of policy evaluations. The computed solution consists in covering the target distribution, up to the necessary connections between the squares and an inevitable density around the initial state.

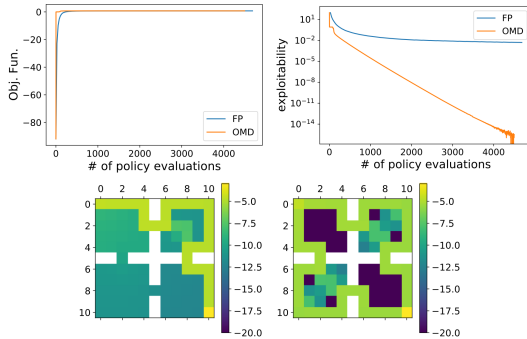


Figure 4: Constrained MDP. Reading order: (a) objective function; (b) exploitability; (c) log-density for the FP population; (d) log-density for the OMD population.

Constrained MDP. Here, the optimized function is $F_{\text{cmdp}}(\mu) = \langle \mu, r \rangle - \frac{\lambda}{2} (\max(0, \langle \mu, c \rangle))^2$, with corresponding reward $\mathcal{R}_{\text{cmdp}}(\cdot, \mu) = \nabla F_{\text{cmdp}}(\mu) = r - \lambda c \max(0, \langle \mu, c \rangle)$. Results are presented Fig. 4, and one can make similar observations regarding FP and OMD (even though the difference of performance is much more noticeable through the exploitability here). There are multiple Nash in this

case. FP favors one mode because best responses are computed using policy iteration (mixture policies are deterministic), while OMD solution is multimodal thanks to the policy being softmax.

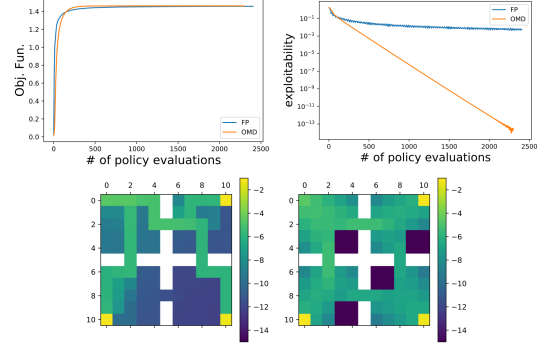


Figure 5: Multi-objective RL. Reading order: (a) objective function; (b) exploitability; (c) log-density for the FP population; (d) log-density for the OMD population.

Multi-objective RL. The function is $F_{\text{morl}}(\mu) = -\sum_{k=1}^3 (1 - \langle \mu, r^k \rangle)^2$ and the reward $\mathcal{R}_{\text{morl}}(\cdot, \mu) = \nabla F_{\text{morl}}(\mu) = 2 \sum_{k=1}^3 r^k (1 - \langle \mu, r^k \rangle)$. Results are presented Fig. 5. Again, similar observations can be made. It might appear less clearly from the log-densities of the final population that OMD is closer to the Nash equilibrium, but it is the case both according to the objective function and to the exploitability (and very clearly here).

Overall, these numerical illustrations suggest that OMD, an approach specifically designed for MFGs, may be a viable alternative to Hazan et al. [21] approach for CURL problems. This highlights an additional evidence of the interest of bridging both communities.

6 CONCLUSION

We have shown that CURL is a subclass of MFGs and linked concepts from both fields: concavity and monotonicity, optimality conditions and exploitability, maximizer and Nash. We also discussed algorithms, showed that Hazan et al. [21] algorithm is indeed FP, and used this connection to FW to provide a rate of convergence for discrete-time FP. Our numerical illustrations suggest that it may be worth considering MFG algorithms for addressing CURL problems.

Solving an MFG from a global planner perspective leads to Mean Field Control (MFC) problems, which have also been considered for interpreting Deep learning algorithms [39]. MFC aims at maximizing $\langle \mu_\pi, \mathcal{R}(\cdot, \mu_\pi) \rangle$ and, as such, may be seen as a special case of CURL, with possible additional difficulties (e.g., population-dependent dynamics or non-concave utility function). Thus, analytical and numerical solutions to MFC (e.g., [31]) may also be useful for CURL. Another interesting perspective would be to study these algorithms in an approximate setting, both theoretically and empirically, something which is still quite preliminary in both fields, compared for example to RL. Lastly, we have seen in Sec. 4.3 that a part of the MFG literature (fixed-point-based) is not readily applicable to CURL. It would be interesting to study if these settings could be reconciled (adapting MFGs algorithms, or considering CURL in a stationary setting).

REFERENCES

- [1] Abbas Abdolmaleki, Jost Tobias Springenberg, Yuval Tassa, Remi Munos, Nicolas Heess, and Martin Riedmiller. 2018. Maximum a Posteriori Policy Optimisation. In *International Conference on Learning Representations*.
- [2] Jacob Abernethy and Jun-Kun Wang. 2017. On Frank-Wolfe and Equilibrium Computation. In *Advances in Neural Information Processing Systems*.
- [3] Mridul Agarwal and Vaneet Aggarwal. 2019. Reinforcement Learning for Joint Optimization of Multiple Rewards. *arXiv preprint arXiv:1909.02940* (2019).
- [4] Berkay Anahtarci, Can Deha Kariksiz, and Naci Saldi. 2019. Fitted Q-learning in mean-field games. *arXiv preprint arXiv:1912.13309* (2019).
- [5] Berkay Anahtarci, Can Deha Kariksiz, and Naci Saldi. 2020. Q-learning in regularized mean-field games. *arXiv preprint arXiv:2003.12151* (2020).
- [6] Sébastien Bubeck. 2015. Convex Optimization: Algorithms and Complexity. *Foundations and Trends® in Machine Learning* 8, 3-4 (2015), 231–357.
- [7] Haoyang Cao, Xin Guo, and Mathieu Laurière. 2020. Connecting GANs, MFGs, and OT. *arXiv e-prints* (2020), arXiv-2002.
- [8] Pierre Cardaliaguet and Saeed Hadikhannloo. 2017. Learning in mean field games: the fictitious play. *ESAIM: Control, Optimisation and Calculus of Variations* 23, 2 (2017), 569–591.
- [9] René Carmona, François Delarue, et al. 2018. *Probabilistic Theory of Mean Field Games with Applications I-II*. Springer.
- [10] Wang Chi Cheung. 2019. Regret Minimization for Reinforcement Learning with Vectorial Feedback and Complex Objectives. In *Advances in Neural Information Processing Systems*.
- [11] Kai Cui and Heinz Koeppl. 2021. Approximately solving mean field games via entropy-regularized deep reinforcement learning. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 1909–1917.
- [12] Luca Di Persio and Matteo Garbelli. 2021. Deep Learning and Mean-Field Games: A Stochastic Optimal Control Perspective. *Symmetry* 13, 1 (2021), 14.
- [13] Romuald Elie, Julien Pérolat, Mathieu Laurière, Matthieu Geist, and Olivier Pietquin. 2020. On the convergence of model free learning in mean field games. In *Proceedings of the AAAI Conference on Artificial Intelligence*. 7143–7150.
- [14] Marguerite Frank and Philip Wolfe. 1956. An algorithm for quadratic programming. *Naval research logistics quarterly* 3, 1-2 (1956), 95–110.
- [15] Matthieu Geist, Bruno Scherrer, and Olivier Pietquin. 2019. A theory of regularized markov decision processes. In *International Conference on Machine Learning*. PMLR, 2160–2169.
- [16] Seyed Kamyar Seyed Ghasemipour, Richard Zemel, and Shixiang Gu. 2020. A divergence minimization perspective on imitation learning methods. In *Conference on Robot Learning (CoRL)*. PMLR, 1259–1277.
- [17] Gauthier Gidel, Tony Jebara, and Simon Lacoste-Julien. 2017. Frank-wolfe algorithms for saddle point problems. In *Artificial Intelligence and Statistics (AISTATS)*. PMLR, 362–371.
- [18] Diogo A Gomes and Vardan K Voskanyan. 2016. Extended deterministic mean-field games. *SIAM Journal on Control and Optimization* 54, 2 (2016), 1030–1055.
- [19] Xin Guo, Anran Hu, Renyuan Xu, and Junzi Zhang. 2019. Learning mean-field games. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [20] Janice H Hammond. 1984. *Solving asymmetric variational inequality problems and systems of equations with generalized nonlinear programming algorithms*. Ph.D. Dissertation. Massachusetts Institute of Technology.
- [21] Elad Hazan, Sham Kakade, Karan Singh, and Abby Van Soest. 2019. Provably efficient maximum entropy exploration. In *International Conference on Machine Learning (ICML)*. PMLR, 2681–2691.
- [22] Minyi Huang, Roland P Malhamé, Peter E Caines, et al. 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.
- [23] Jean-Michel Lasry and Pierre-Louis Lions. 2007. Mean field games. *Japanese journal of mathematics* 2, 1 (2007), 229–260.
- [24] Lisa Lee, Benjamin Eysenbach, Emilio Parisotto, Eric Xing, Sergey Levine, and Ruslan Salakhutdinov. 2019. Efficient exploration via state marginal matching. *arXiv preprint arXiv:1906.05274* (2019).
- [25] Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. 2020. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643* (2020).
- [26] David Mguni, Joel Jennings, and Enrique Munoz de Cote. 2018. Decentralised learning in systems with many, many strategic agents. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32.
- [27] Yurii Nesterov. 1998. *Introductory lectures on convex programming volume i: Basic course*. Springer.
- [28] George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. 2021. Normalizing flows for probabilistic modeling and inference. *Journal of Machine Learning Research* 22, 57 (2021), 1–64.
- [29] Julien Perolat, Sarah Perrin, Romuald Elie, Mathieu Laurière, Georgios Piliouras, Matthieu Geist, Karl Tuyls, and Olivier Pietquin. 2022. Scaling up Mean Field Games with Online Mirror Descent. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.
- [30] Sarah Perrin, Julien Pérolat, Mathieu Laurière, Matthieu Geist, Romuald Elie, and Olivier Pietquin. 2020. Fictitious play for mean field games: Continuous time analysis and applications. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [31] Laurent Pfeiffer. 2017. Numerical Methods for Mean-Field-Type Optimal Control Problems. *arXiv preprint arXiv:1703.10001* (2017).
- [32] Martin L Puterman. 2014. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons.
- [33] Julia Robinson. 1951. An iterative method of solving a game. *Annals of mathematics* (1951), 296–301.
- [34] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. 2015. Trust region policy optimization. In *International conference on machine learning*. PMLR, 1889–1897.
- [35] Jayakumar Subramanian and Aditya Mahajan. 2019. Reinforcement learning in stationary mean-field games. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*. 251–259.
- [36] Ahmed Touati, Amy Zhang, Joelle Pineau, and Pascal Vincent. 2020. Stable policy optimization via off-policy divergence regularization. In *Conference on Uncertainty in Artificial Intelligence*. PMLR, 1328–1337.
- [37] Nino Vieillard, Tadashi Kozuno, Bruno Scherrer, Olivier Pietquin, Rémi Munos, and Matthieu Geist. 2020. Leverage the average: an analysis of KL regularization in reinforcement learning. In *Advances in Neural Information Processing Systems*.
- [38] Nino Vieillard, Olivier Pietquin, and Matthieu Geist. 2020. Munchausen reinforcement learning. In *Advances in Neural Information Processing Systems*.
- [39] E Weinan, Jiequn Han, and Qianxiao Li. 2019. A mean-field optimal control formulation of deep learning. *Research in the Mathematical Sciences* 6, 1 (2019), 1–41.
- [40] Qiaomin Xie, Zhuoran Yang, Zhaoran Wang, and Andreea Minca. 2020. Provable fictitious play for general mean-field games. *arXiv preprint arXiv:2010.04211* (2020).
- [41] Jiachen Yang, Xiaojing Ye, Rakshit Trivedi, Huan Xu, and Hongyuan Zha. 2018. Deep mean field games for learning optimal behavior policy of large populations. In *International Conference on Learning Representations*.
- [42] Yaodong Yang, Rui Luo, Minne Li, Ming Zhou, Weinan Zhang, and Jun Wang. 2018. Mean field multi-agent reinforcement learning. In *International Conference on Machine Learning*. PMLR, 5571–5580.
- [43] Tom Zahavy, Brendan O’Donoghue, Guillaume Desjardins, and Satinder Singh. 2021. Reward is enough for convex MDPs. In *Advances in Neural Information Processing Systems*.
- [44] Junyu Zhang, Alec Koppel, Amrit Singh Bedi, Csaba Szepesvari, and Mengdi Wang. 2020. Variational policy gradient method for reinforcement learning with general utilities. In *Advances in Neural Information Processing Systems (NeurIPS)*.