

Foolproof Cooperative Learning

Alexis Jacq

Google Research, Brain Team

ALEXISJACQ@GOOGLE.COM

Julien Perolat

DeepMind

PEROLAT@GOOGLE.COM

Matthieu Geist

Google Research, Brain Team

MFGEIST@GOOGLE.COM

Olivier Pietquin

Google Research, Brain Team

PIETQUIN@GOOGLE.COM

Editors: Sinno Jialin Pan and Masashi Sugiyama

Abstract

This paper extends the notion of learning algorithms and learning equilibriums from repeated games theory to stochastic games. We introduce Foolproof Cooperative Learning (FCL), an algorithm that converges to an equilibrium strategy that allows cooperative strategies in self-play setting while being not exploitable by selfish learners. By construction, FCL is a learning equilibrium for repeated symmetric games. We illustrate the behavior of FCL on symmetric matrix and grid games, and its robustness to selfish learners.

1. Introduction

Learning equilibriums describe a set of learning algorithms used by a team of agents to learn repeated games such that no agent can modify its algorithm and increase its average payoff (Brafman and Tennenholtz, 2003). In that setting, a learning algorithm is viewed as a strategy that requires no initial information about the played game, so it works with any game of a given set.

Besides, just like the Folk theorem allows to construct cooperative Nash equilibriums for repeated games based on retaliation strategies, it is possible to construct a cooperative learning equilibrium. Therefore, a team discovering a game involving a social dilemma can agree on a set of learning rules that prevent an eventual selfish player to learn a defecting behavior. In this paper, we build a cooperative learning equilibrium that applies to the more general set of stochastic games (Shapley, 1953) involving sequential decision making.

Multi-agent reinforcement learning (MARL) brings a framework to construct algorithms that aim to solve stochastic games where players individually or jointly search for an optimal decision-making policy to maximize a reward function. Individualist approaches mostly aim at reaching Nash equilibriums, taking the best actions whatever the opponents behaviors are (Bowling and Veloso, 2001; Littman, 2001). Cooperative approaches aim at optimizing a common objective. If agents can agree on joint actions, cooperation can be viewed as a single agent problem in a larger dimension (Claus and Boutilier, 1998). However, when rewards are individual but agents agree to learn a joint cooperative approach, they remain

easily exploited if one agent starts being individualist. In that case, a team needs a learning equilibrium that safely leads rational agents to learn the common objective rather than the maximization of their own rewards.

We focus on symmetric situations where no agent has an individual advantage. For example, this is the case on an island with a quantity of resources equally accessible to all agents. Moreover, we consider repeated stochastic games (Munoz de Cote and Littman, 2008), modelling the situation where players restart from the beginning after reaching absorbing states or after a finite number of time steps. In the island resource example, repetitions could represent successive days or 4-seasons cycles. In fact, any repeated stochastic game where all players have the same reward function and are subject to the same transition laws is symmetric. This applies to most of common-pool resource appropriation games (Pérolat et al., 2017) and other sequential social dilemmas (Leibo et al., 2017).

In this context, we extend the definition of learning algorithms – mapping between histories and probability distributions – and the definition of learning equilibriums to stochastic games. We introduce Foolproof Cooperative Learning (FCL), an algorithm that both learns cooperative and retaliating strategies. In that perspective, one FCL player i learns two value-functions for each other player j , the one estimating the j 's payoff for cooperating when all the team cooperates, and the other estimating j 's best possible payoff when all the team tries to lower its score in an eventual retaliation. When a player deviates from the agreed exploring behavior and is not cooperative during greedy steps, the rest of the team retaliates in consequence. We show that FCL is a learning equilibrium forcing a cooperative behavior, and we empirically verify this claim with two-agent matrix games and grid-world repeated symmetric games. Our contributions are (1) a theoretical framework and an experimental setting to study learning equilibriums for stochastic games and (2) the construction of a cooperative learning equilibrium using finite retaliations and so allowing selfish learners to eventually cooperate after exploring defections.

2. Definitions and Notations

An N -player stochastic game can be written as a tuple $(\mathcal{S}, (\mathcal{A}_i)_{i=1\dots N}, \mathcal{P}, \mu_0, (r_i)_{i=1\dots N})$, where \mathcal{S} is the set of states, \mathcal{A}_i the set of actions for player i , \mathcal{P} the transition probability $(\mathcal{P}(\cdot|s, a_1 \dots a_N))$, μ_0 a distribution over initial states $(\mu(s^0))$, r_i the reward function for player i $(r_i(s, a_1 \dots a_N))$. We also assume bounded, deterministic reward functions and finite state and action spaces. In a repeated stochastic game, a stochastic game (the stage game) is played and terminates when at least one player reaches an absorbing state, or after a finite number of steps. This is repeated an infinite number of times, and players have to maximize their average return during a stage game (Munoz de Cote and Littman, 2008).

A stationary strategy for player i , $\pi_i(\cdot|s) \in \Pi_{\mathcal{A}_i}$, maps a state to a probability distribution over its set of possible actions. We write π_{-i} the product of all players strategies but player i and $\boldsymbol{\pi} = \pi_1 \times \dots \times \pi_N = \pi_i \times \pi_{-i}$ the product of all players strategies (the strategy profile). Given opponents strategies π_{-i} , the goal for a rational player i is to find a strategy π_i^* that maximizes its average payoff \mathcal{R}_i during a stage game:

$$\pi_i^* = \operatorname{argmax}_{\pi_i} \mathcal{R}_i(\pi_i, \pi_{-i}) = \operatorname{argmax}_{\pi_i} \mathbb{E}_{\pi_i, \pi_{-i}, \mathcal{P}} \sum_l \gamma^l r(s^l, a_i^l, a_{-i}^l).$$

The policy π_i^* depends on the opponents strategies and is called the best response for player i to π_{-i} . In general, we call strategy any process $\{\pi^t\}_t$ defining a stationary strategy for any stage t . The value of a player's non-stationary strategy $\{\pi^t\}_t$ is the average return over stage games, $\mathbb{E}_{t>0} [\mathcal{R}_i(\pi_i^t, \pi_{-i}^t)] = \liminf_{T \rightarrow \infty} \mathbb{E} \left[\sum_{t=0}^T \mathcal{R}_i(\pi_i^t, \pi_{-i}^t) \right]$. A Nash equilibrium describes a stationary strategy profile $\pi^* = \pi_1^* \times \dots \times \pi_N^*$, such that no player can individually deviate and increase its payoff (Nash, 1951):

$$\forall 1 \leq i \leq N, \forall \pi_i \in \Pi_{\mathcal{A}_i}, \mathcal{R}_i(\pi_i, \pi_{-i}^*) \leq \mathcal{R}_i(\pi_i^*, \pi_{-i}^*).$$

This definition can be extended to non-stationary strategies using expected return over stage games: no player can individually deviate from an equilibrium non-stationary strategy and increase its average payoff:

$$\forall 1 \leq i \leq N, \forall \{\pi_i^t \in \Pi_{\mathcal{A}_i}\}_t, \mathbb{E}_{t>0} [\mathcal{R}_i(\pi_i^t, \pi_{-i}^{t,*})] \leq \mathbb{E}_{t>0} [\mathcal{R}_i(\pi_i^{t,*}, \pi_{-i}^{t,*})].$$

As $\mathbb{E}_{t>0} [\mathcal{R}_i(\pi_i^t, \pi_{-i}^t)] = \mathcal{R}_i(\pi_i, \pi_{-i})$ for stationary strategy profiles, any stationary strategy equilibrium is still an equilibrium among non-stationary processes.

In order to allow rewarding or retaliation strategies, we only consider games where all players are aware of all opponents actions and rewards, and receive a signal each time the game is reset. We also admit players to share information with some opponents in order to organize joint retaliation actions or joint explorations. Moreover, we only consider *Repeated Symmetric Games* (RSG):

Definition 1 (Repeated Symmetric game (RSG)) *A repeated stochastic game is symmetric if, for any stationary strategy profile $(\pi_1 \dots \pi_N)$ and any permutation ψ over players:*

$$\forall 1 \leq i \leq N, \mathcal{R}_{\psi(i)}(\pi_i, \pi_{-i}) = \mathcal{R}_i(\pi_{\psi(i)}, \pi_{\psi(-i)}).$$

This generalizes the definition for symmetric N -player matrix games (Dasgupta et al., 1986) to stochastic games where players' utilities are replaced by average returns¹. Note that in a symmetric game, for any Nash equilibrium with returns $(\mathcal{R}_1 \dots \mathcal{R}_N)$ and for any permutation σ over players, there is another Nash equilibrium with returns $(\mathcal{R}_{\sigma(1)} \dots \mathcal{R}_{\sigma(N)})$. In this paper, we use the concept of N -cyclic permutations to construct specific strategies:

Definition 2 (N -cyclic permutation) *A permutation σ is N -cyclic if for all $i, j \in \{1 \dots N\}$, there is k such that $\sigma^k(i) = j$.*

3. Cooperation and retaliation

We say a strategy is cooperative if it maximizes a common quantity $\hat{\mathcal{R}} = f(\mathcal{R}_1 \dots \mathcal{R}_N)$. Usual examples are strategies that maximize the sum, the product or the minimum of players returns. In RSGs, the strategy that maximizes the minimum of player returns is particularly interesting as it coincides with the best egalitarian payoff. In this paper, we refer to this strategy as the *egalitarian strategy*. An important property of RSGs is the fact that egalitarian strategies can always be obtained by repeatedly applying an N -cyclic permutation on a stationary strategy that maximizes the sum of players returns (see appendix B):

1. Actually, the definition of Dasgupta et al. (1986), $\forall i, \mathcal{R}_i(\pi_i, \pi_{-i}) = \mathcal{R}_{\psi(i)}(\pi_{\psi(i)}, \pi_{\psi(-i)})$, is incorrect: symmetries are not independent of player identities, which is not the case if the right-hand return is indexed with the inverse permutation instead (Vester, 2012).

Theorem 1 Let π_i^Σ be a stationary strategy for player i that maximizes the sum of players returns in an N -player RSG, σ an N -cyclic permutation over players, and t indexing the successive stages. Then, the strategy $\pi^t = (\pi_{\sigma^t(1)}^\Sigma \dots \pi_{\sigma^t(N)}^\Sigma)$ (where $\sigma^t = \sigma \circ \dots \circ \sigma$ t times) is an egalitarian strategy.

In repeated games, players can retaliate when a selfish one deviates from a target strategy. In that case, the target strategy is said to be enforceable (Osborne and Rubinstein, 1994): if all players are accorded to retaliate when one player deviates from a strategy profile, no player can improve its payoff by individually deviating from the strategy profile. This phenomenon is described by the Folk theorem and allows building equilibrium strategies achieving cooperative goals. In an RSG, the egalitarian strategy is enforceable since for any player i , its payoff V_i^c for cooperating is always larger than its minimax payoff V_i^r when all other players retaliate to lower its reward (because of the symmetric and convex shape of the space of possible payoffs (Munoz de Cote and Littman, 2008)). When a single retaliation is too small so it is still worth defecting for a selfish player, it must be repeated. As Littman and Stone (2005), one can compute a minimal number of retaliations to enforce the cooperation. Let V_i^d be the payoff obtained by player i for defecting, then

$$K_j = \left\lceil \frac{V_j^d - V_j^c}{V_j^c - V_j^r} \right\rceil \tag{1}$$

is a sufficient number of retaliations (see Appendix A). In the edge case where $V_j^c = V_j^r$, the required number of retaliations becomes infinite, but the cooperative objective is not affected (this is the case in rock–paper–scissors). Let π_{Folk} be the (non-stationary) strategy that follows π^* if all players cooperate, or repeat a minimax retaliation over K_j stage games if a player j deviates from π^* . By construction, π_{Folk} is a Nash equilibrium.

4. Learning algorithm

In normal form games, a *learning algorithm* maps an history of actions and rewards to a probability distribution (Brafman and Tennenholtz, 2003). In stochastic games, one must consider a mapping between an history of transitions during stage games, $\mathcal{H}^T = \{\{s^l, a_i^l, a_{-i}^l, r_i^l, r_{-i}^l\}_{l \in t}\}_{t < T}$, to state-dependent strategies $\pi(\cdot|s)$. For simplification, we will write² $A_i(T) = A_i(\mathcal{H}^T) = \pi_i^T$ the algorithm used by a player i . The algorithm profile $\mathbf{A} = (A_1 \dots A_N)$ is the set of all players algorithms.

4.1. Multi-agent learning

Reinforcement learning provides a class of algorithms that aim at maximizing an agent’s return. Out of all of them, our interest concerns Q -learning approaches (Watkins and Dayan, 1992) for three reasons: they are model-free, off-policy and they are guaranteed to converge in finite state and action spaces. In a game \mathcal{G} , for a player i and given opponents policy π_{-i} , the basic idea is to learn a Q -function that approximates, for each state and action, the

2. Brafman and Tennenholtz (2003) directly treated learning algorithms like strategies (considering the meta game of choosing an algorithm for a set of games). In this paper, we preferred to keep a distinct notation in order to stay consistent with the reinforcement learning notation.

average return starting from playing that specific action at that given state while using the best strategy afterwards. For the Q -function Q_i , associated with player i 's optimal policy (the one that maximizes its return), the following holds:

$$Q_i(s, a_i, a_{-i}) = r_i(s, a_i, a_{-i}) + \gamma \sum_{s'} \mathcal{P}(s'|s, a_i, a_{-i}) \max_{a'_i} Z_i(a'_i, s', \pi_{-i}),$$

where $Z_i(a_i, s, \pi_{-i}) = \sum_{a'_{-i}} \pi_{-i}(a'_{-i}|s') Q_i(s', a'_i, a'_{-i})$ is the expected value for agent i given its opponents policy. The Q -learning algorithm is constructed in order to progressively approximate the Q -function without knowledge of the problem dynamics \mathcal{P} and reward functions r , and without knowing the decision process that generated the history buffer (in contrast, for example, to policy gradient algorithms (Williams, 1992)). In finite state and action spaces, the approximation is obtained by successively applying the updates:

$$Q_i^{t+1}(s^t, a_i^t, a_{-i}^t) = Q_i^t(s^t, a_i^t, a_{-i}^t) + \alpha_t (r_i^t + \gamma \max_{a_i} Z_i(a_i, s^{t+1}, \pi_{-i}) - Q_i^t(s^t, a_i^t, a_{-i}^t)),$$

where α_t is the learning rate. However, when the opponent policy is not fixed, maximizing the Q -function with respect to actions is no longer an improvement of the policy (the response of the opponents to this deterministic policy can decrease the average player's return). MARL provides several alternative greedy improvements. For example, a defensive player can expect opponents to minimize its Q -function (*minimax* Q -learning). In that case, a greedy improvement of the policy to evaluate the value of a new state is obtained by solving the linear problem (Littman, 1994):

$$\pi_i^{greedy}(\cdot|s) = \operatorname{argmax}_{\pi_i} \min_{a_{-i}} \sum_{a_i} \pi_i(a_i|s) Q_i(s, a_i, a_{-i}) = \operatorname{argmax}_{\pi_i} \min_{a_{-i}} Z_{-i}(a_{-i}, s, \pi_i)$$

and the corresponding Q -learning update becomes:

$$Q_i^{t+1}(s^t, a_i^t, a_{-i}^t) = Q_i^t(s^t, a_i^t, a_{-i}^t) + \alpha_t (r_i^t + \gamma \max_{\pi_i} \min_{a_{-i}} Z_{-i}(a_{-i}, s, \pi_i) - Q_i^t(s^t, a_i^t, a_{-i}^t)).$$

Besides, even when the game is cooperative, the fact that all strategies are changed at each update breaks the Markov hypothesis. If agents cannot communicate, it becomes complex to understand how one agent personally contributed to improving or deteriorating the strategy profile (Foerster et al., 2018b). This is not the case in this paper, as we consider the case where agents can fully communicate. Thus, learning a cooperative joint strategy becomes equivalent to solving a larger dimensional MDP and the optimal (cooperative) strategy is iteratively computed:

$$Q^{c,t+1}(s^t, a_i^t, a_{-i}^t) = Q^{c,t}(s^t, a_i^t, a_{-i}^t) + \alpha_t (\sum_i r_i^t + \gamma \max Q^{c,t}(s^{t+1}, \cdot) - Q^{c,t}(s^t, a_i^t, a_{-i}^t)). \quad (2)$$

In order to learn a retaliating behavior with a finite number of retaliations, a team of agents must learn a defensive strategies against a single defecting agent (j), which corresponds to the opposite of *minimax* Q -learning indexing where one agent is defensive against others:

$$Q_j^{r,t+1}(s^t, a_i^t, a_{-i}^t) = Q_j^{r,t}(s^t, a_i^t, a_{-i}^t) + \alpha_t (r_i^t + \gamma \max_{a_j} \min_{\pi_{-j}} Z_j(a_j, s, \pi_{-j}) - Q_j^{r,t}(s^t, a_i^t, a_{-i}^t)). \quad (3)$$

For estimating a finite number of retaliations with Eq. (1), agents must learn the maximum benefit one agent can obtain by defecting while others are cooperating, V_j^d . Knowing Q^c , this value can also be derived from Q-learning:

$$Q_j^{d,t+1}(s^t, a_i^t, a_{-i}^t) = Q_j^{d,t}(s^t, a_i^t, a_{-i}^t) + \alpha_t \left(\sum_i r_i^t + \gamma V_j^d(s^{t+1}) - Q_j^{d,t}(s^t, a_i^t, a_{-i}^t) \right) \quad (4)$$

with $V_j^d(s) = \max_{a'_j} Q_j^d(s, a'_j, \operatorname{argmax}_{a_{-j}} \max_{a_j} Q^c(s, a_j, a_{-j}))$.

4.2. Learning equilibrium

Learning equilibriums are defined over normal form games (Brafman and Tennenholtz, 2003), but one could consider any stochastic game as a normal form game where possible decisions are the set of all possible strategies $\pi(\cdot|s) \in \Pi_{\mathcal{A}}$. In that sense, the original definition applies to stochastic games:

Definition 3 (Learning equilibrium) *Let \mathcal{G} be a set of stochastic games. An algorithm profile $\mathbf{A}^* = (A_1^* \dots A_N^*)$ is a learning equilibrium for \mathcal{G} if, for any game $g \in \mathcal{G}$, there is a stage T_g such that, for any player i and any learning algorithm A_i :*

$$\mathbb{E}_{t>T_g} \mathcal{R}_i(A_i(t), A_{-i}^*(t)) \leq \mathbb{E}_{t>T_g} \mathcal{R}_i(A_i^*(t), A_{-i}^*(t)).$$

Just like for Nash equilibriums, no player can individually follow an alternative algorithm and increase its asymptotic score. However, one important difference is the fact that a learning algorithm is not defined with respect to a particular game, but a set of games. We may think that a process always playing a Nash equilibrium of the given game ($\pi_i^t = \pi_i^*$ for all t) is a learning equilibrium. However, such a process requires an initial knowledge of the game and can't be obtained from a process starting with an empty condition. Therefore, it can't be described as a learning algorithm. For the same reason, the strategy π_{Folk} described in Sec. 3 is not a learning equilibrium. But one can construct a learning equilibrium that learns to play π_{Folk} : this is the idea of FCL.

5. Foolproof cooperative learning

FCL, as described in Alg. 1 (for a player i), converges to π_{Folk} over the explored set of states. In an N -player game, each FCL player approximates $2N + 1$ Q -functions: one associated with the cooperative policy that maximizes the sum of all players (Q^c), N associated with retaliation policies preventing any defection from other players j (Q_j^r), and N associated with each opponent's best response to the cooperative strategy (learned through Q_j^d with Eq.(4)). At each stage game, FCL will play according to an egalitarian strategy (learned through Q^c with Eq.(2) and applying an N -cyclic permutations to the associated strategy profile) unless one of the opponents deviates from that strategy. In case of a defection, all FCL agents will agree on a joint retaliation based on a defensive strategy (learned through Q_j^r with Eq.(3)) for K stages (according to Eq.(1)). In order to allow exploration, a deterministic process is used to decide, at each time t , between exploration and exploitation. We highlight that FCL differs from pure *minimax* Q-learning in the fact that it also learns to cooperate, and differs

Algorithm 1 FCL for player i .

Input: initial functions Q^c , $\{Q_i^d\}_{i=1\dots N}$ and $\{Q_i^r\}_{i=1\dots N}$, and N-cyclic permutation σ .**For** stages $t = 1$ **to** $+\infty$: **While** stage continue **If** no need for retaliation If exploration time: **explore** with $a_i \sim \text{uniform}$ Otherwise: **cooperate** with $a_i \sim \pi_{\sigma^t(i)}^c(Q^c)$ **Else** (player j previously defected) **retaliate** with $a_i \sim \pi_{i,j}^r(Q_j^r)$ Observe a_{-i} and new state s' , receive reward r_i and observe r_{-i} Update Q^c according to (2) **For** all other agents $j \neq i$ Update Q_j^r according to (3) Update Q_j^d according to (4) **If** it was not an exploration and j did not cooperate

Compute the number of required retaliations according to (1)

from Friend-or-Foe Q -learning (FFQ) (Littman, 2001) which learns to play cooperatively if an opponent is cooperative, or defensively if the opponent is defective with a single Q -function, while FCL learns both cooperation and defense in a disentangled way. Thm 2 states that FCL is a learning equilibrium for RSGs forcing a cooperative behavior.

Theorem 2 *FCL is a learning equilibrium for RSGs.*

In the proof of Thm 2 (see Appx. C), we show that any team of communicating FCL agents converges to π_{Folk} . As constructed, FCL can be seen as a team rule designed to prevent the deviation of one agent of the team. This assumes that agents can be centralized (in order to communicate the times of exploration) and that no sub-team of agent can simultaneously start to deviate and defect. In the case where agents are unable to communicate, they need to estimate if a deviation is caused by an exploration or is resulted by the exploitation of a selfish behaviour. Such an ability, studied in the context of simple matrix games (Ashlagi et al., 2006), would likely improve the robustness of FCL which, so far, requires a relatively high level of coordination. If more than one agent simultaneously start to defect, this sub-team could be perceived as one “meta-opponent” to punish during a longer time. This opens another perspective of future improvement.

6. Experiments

We first explored the case of well known repeated symmetric matrix games involving two players: Iterated Prisoners Dilemma (IPD), Iterated Chicken (ICH) and Rock-Paper-Scissors (RPS). We also studied an altruist version of the IPD (aIPD), where the optimal way to share the outcomes in average is when the two players alternate between (player 1 cooperate, player 2 defect) with rewards (-3, 10) and (player 1 defect, player 2 cooperate), with rewards (10, -3): in average, both players get 3.5. Table 1 shows the payoff matrices. Then, we investigated larger state spaces with grid games known to induce coordination problems and

Table 1: Payoff matrices used for IPD, aIPD, ICH and RPS.

IPD		
	C	D
C	(-1,-1)	(-3,0)
D	(0,-3)	(-2,-2)

aIPD		
	C	D
C	(-1,-1)	(-3,10)
D	(10,-3)	(-2,-2)

ICH		
	Swerve	Straight
Swerve	(2,2)	(1,3)
Straight	(3,1)	(0,0)

RPS			
	R	P	S
R	(0,0)	(-1,1)	(1,-1)
p	(1,-1)	(0,0)	(-1,1)
S	(-1,1)	(1,-1)	(0,0)

Table 2: Grid games. A is the starting position of one player, B is the starting position of the other. At each turn, both players simultaneously select one action among going up, down, left, right or stay. When reward cells with \$ symbol are reached by one player, the player obtains the corresponding reward and the game is immediately reset. $\$_{A:X}$ means that only player A gets the reward X when reaching the cell, $\$_X$ means that any player gets reward X when reaching the cell, and $\$_{X,Y}$ means that the player who reaches the cell gets X and the other gets Y (if the other player reach another rewarding cell, the rewards are summed). Two players can not be on the same cell at the same time and they can not cross each other. In case of conflict, one player reaches the cell and the other stays with probability 0.5. Grey cells are walls and are not reachable.

(a) Grid prisoners dilemma									
					\$100				
\$ _{A:100}			A		B				\$ _{B:100}

(b) Compromise									
					\$ _{B:100}				
			A						\$ _{A:100}

(c) Coordination			
\$ _{B:100}			\$ _{A:100}
	A		B

(d) Temptation			
\$ _{20,-10}	A	B	\$ _{20,-10}
\$ _{40,-20}			\$ _{40,-20}

social dilemmas (Munoz de Cote and Littman, 2008). We introduced a new grid game, closer to the concept of limited resource appropriation: the Temptation game. In Temptation, making a movement to the sides can be seen as taking immediately the resource, while making a movement to the bottom can be seen as waiting for the winter. All grid games are described in details in Table 2. In order to verify that FCL is a learning equilibrium, we compare the score obtained by FCL and by selfish learning algorithms, Q -learning and policy-gradient (PG), against FCL. Indeed, we expect a learning equilibrium performing better than any other algorithm when the opponents are following the learning equilibrium: knowing the other players are using FCL, the rational choice is to use FCL. The reason why we did not compare FCL to more intelligent players was to illustrate its robustness to defectors, in contrast to less naive agents which are easier to cooperate with (like FFQ).




Figure 1: Matrix games. Average scores over 20 runs obtained by two standard RL algorithms and FCL, playing against FCL. In IPD and ICH, after some iterations selfish behaviors, as induced by Q-learning and PG, start being sub-optimal because of FCL retaliations and accumulate less return than a cooperative behaviors, as induced by FCL against itself. In RPS, FCL learns to play with a uniform distribution against selfish algorithms so their average score is null. Black dotted line represents the average score after convergence of two selfish agents playing against themselves (the *minimax* solution).

The robustness of FFQ to selfish players could be compared with the robustness of FCL to selfish players: they would be similar, the difference being the fact that FCL can start cooperating if the opponent stops defecting, while FFQ only converges (by construction) to a defensive behaviour when facing a selfish player.

6.1. Implementation details

We implemented FCL using a state-dependent learning rate $\alpha_t = (\sum_{l < t} \delta\{s^l = s^t\})^{-1}$ that counts the number of state visits, and explorations are determined by a process $\phi(t) = \{X_t > \epsilon d^t\}$ where X_t is a pseudo-random uniform sample between 0 and 1 with a communicated seed, ϵ a threshold and d a decay parameter. For selfish Q-learning, we used a similar learning rate and exploration process, however with different seeds and decay parameters. The policy gradient was implemented with a tabular representation and Adam gradient descent with learning rate 0.1. Since matrix games are not sequential and since grid games were automatically reset after 30 steps, we could use a discount factor $\gamma = 1$ to estimate value functions. In practice, we found that adding 1 to the minimal number of retaliation repeats given in Eq. 1 significantly improves the robustness to selfish learners. In iterated matrix games, since they do not require large explorations, we used $\epsilon = 0.5$ and $d = 0.9$ for both selfish Q-learning and FCL. We used $\epsilon = 1$ and $d = 0.995$ in grid games.

6.2. Results

In order to experimentally verify that FCL behaves as a learning equilibrium, we observe the score obtained by its opponent: if the opponent cooperates (for example if the opponent is itself an FCL), it should always obtain a better score than by following any other behavior (if the opponent follows a selfish policy gradient or a Q-learning algorithm). Therefore we compare the score of an FCL player with the scores of policy-gradient and Q-learning players, while playing with an FCL opponent. Figure 1 displays our results with the three




Figure 2: Grid games. Average scores over 20 runs obtained by two standard RL algorithms and FCL, playing against FCL. After some iterations, selfish behaviors, as induced by Q-learning and PG, start being sub-optimal because of FCL retaliations and accumulate less return than a cooperative behavior, as induced by FCL against itself. Black dotted line represents the average score after convergence of two selfish agents playing against themselves (the *minimax* solution).

matrix games IPD, aIPD, ICH and RPS. Figure 2 displays our results on grid games. As expected, the score of selfish learners was never higher than the score of FCL, when the opponent is FCL. Except in RPS, defection conduced to less rewards than cooperation because of retaliations. In RPS, FCL found that the only way to retaliate was by infinitely playing randomly against selfish learners, resulting in an average of 0 reward for all players, equivalent to the reward for cooperation. This illustrates the fact that FCL is a learning equilibrium, since no tested algorithm performs better than FCL against FCL. Consequently, FCL was never exploited by selfish learners while being cooperative in self-play.

6.3. More than 2 players

We also investigated the robustness of FCL when more than 2 players are involved. What is expected is that the best a single agent can do when facing a team of FCL is to play as FCL. In order to verify this clam, we created a social dilemma game involving $N > 2$ players. It is a tensor game where N agents are given a cake, and may choose between sharing, robing or poisoning. If no player robs, each sharing player receives a reward $\frac{1-c}{n_s}$, where n_s is the number of sharing agents and $c = -\frac{n_p}{(N-1)}$ the eventual cost from poison (where n_p is the number of poisoning agents). If n_r agents decide to rob, one of them receives a reward $0.5 - c$ with probability $\frac{1}{n_r}$, all others receive nothing. Poisoning agents receive no reward. The

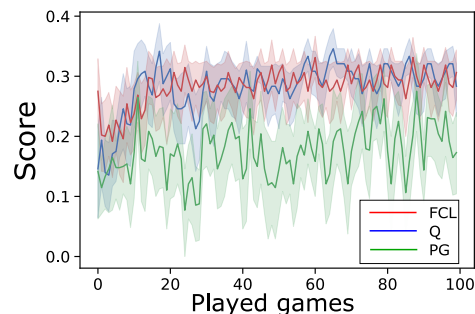


Figure 3: Grid games. Average scores over 20 runs obtained by two standard RL algorithms and FCL, playing against a team of 2 FCL. Q-learning players learned to cooperate with the team of FCL, leading to the same score than the one obtained by three FCL players ($1/3$ if all players chose to share).

cost from poison is designed to equal one if all but one agent poisoned the cake while the last one decided to rob it. If all agents cooperate, the best they can do is to share the cake, obtaining a reward of $\frac{1}{N}$ (0.333 if $N = 3$). But if one single agent starts to rob the cake (while all other are sharing), it obtains a reward of 0.5. Figure 3 shows the results when a team of two FCL are facing a Q-learning, a PG or a third FCL. It appears that Q-learning players learned to cooperate with the team of FCL, leading to the same score than the one obtained by three FCL players.

7. Related work

Learning cooperative behaviors in a multi-agent setting is a vast field of research, and various approaches depend on assumptions about the type of games, the type and number of agents, the type of cooperation and the initial knowledge.

When the game’s dynamics are initially known and in two-player settings, an egalitarian strategy can be obtained by mixing dynamic and linear programming. Therefore, a polynomial-time algorithm can be used to solve repeated matrix games (Littman and Stone, 2005), as well as repeated stochastic games (Munoz de Cote and Littman, 2008). Since an egalitarian strategy is always better than a *minimax* strategy (the disagreement point) (Osborne and Rubinstein, 1994), a cooperative Nash equilibrium is immediately given. An alternative to our cooperate or retaliate architecture consists in choosing between maximizing oneself reward (being competitive) or maximizing a cooperative reward, for example by inferring opponents intentions (Kleiman-Weiner et al., 2016). In contrast, FCL does not require the model in order to construct a foolproof cooperative behavior.

In games inducing social dilemmas and when the dynamics are accessible as an oracle, cooperative solutions can also be obtained by self-play and then applied to define a retaliating behavior forcing cooperation (Lerer and Peysakhovich, 2017), even when opponent actions are unknown, since in that case the reward function already brings sufficient information (Peysakhovich and Lerer, 2018). Here again, they use an offline procedure which does not apply to our purely online setting.

Closer to our setting, when the dynamics are unknown, online MARL can extract cooperative solutions in some non-cooperative games, and particularly in restricted resource appropriation (Pérolat et al., 2017). Using alternative objectives based on all players reward functions and their propensity to cooperate or defect improves and generalizes the emergence of cooperation in non-cooperative games and limits the risk of being exploited by purely selfish agents (Hughes et al., 2018). Like FFQ (Littman, 2001), the novelty of FCL is to disentangle the cooperative and the retaliating strategies so it can always switch from one behavior to the opposite without a forgetting and re-learning phase.

A similar approach, called Learning with Opponent Learning Awareness (LOLA), consists in modelling the strategies and the learning dynamics of opponents as part of the environment’s dynamics and to derive the gradient of the average return’s expectation (Foerster et al., 2018a). If LOLA has no guaranties of convergence, a recent improvement of the gradient computation, which interpolates between first and second-order derivations, is proved to converge to a local optimums (Letcher et al., 2018). Although such agents are purely selfish, empirical results show that they are able to shape each others learning trajectories and to cooperate in prisoners dilemma. A limitation of this approach towards

building learning equilibrium is the strong assumption regarding the opponents learning algorithms, supposed to perform policy gradient. Also, this approach differs to our goal since LOLA is selfish and aims at shaping an opponent’s behavior (in 2-player settings) while FCL is cooperative but retaliates in response to selfish agents (in N -player settings).

Learning equilibrium solutions have been constructed for repeated matrix games (Brafman and Tennenholtz, 2003; Ashlagi et al., 2006). However, these solutions would not easily adapt to stochastic games, one main reason being the fact that exploration becomes infinite, while it only requires $A \times N$ steps in N -agent matrix games with A different actions. Consequently, after a finished phase of exploration in matrix games, the deterministic payoff matrix is known and they can extract a Nash Equilibrium to exploit. Note that the restriction to symmetric games seems recurrent in learning equilibrium literature (Brafman and Tennenholtz, 2005; Tennenholtz and Zohar, 2009). In repeated congestion games, it is even possible to construct a class of asymmetric games that does not admit any learning equilibrium, hence the importance of the symmetry assumption.

8. Conclusion

We introduced FCL, a model-free learning algorithm that, by construction, converges to an equilibrium strategy, cooperative with itself and retaliating when selfish algorithms are defecting. We proposed a definition for learning equilibrium, describing a class of learning algorithms such that the best way to play against it is to adopt the same behavior. We demonstrated that FCL is a learning equilibrium that forces a cooperative behavior, and we empirically verified this claim in various settings. Our approach could be improved by facilitating opponent’s learning of the optimal cooperative response and by using faster learning approaches. It could also be adapted to larger dimensions such as continuous state spaces and partially observed settings with function approximation by replacing tabular Q -learning with deep Q -learning (Mnih et al., 2015).

Appendix A. Sufficient number of retaliations.

Proof Since $K_j \geq \frac{V_j^d - V_j^c}{V_j^c - V_j^r}$ and $V_j^c \geq V_j^r$, we write:

$$K_j(V_j^c - V_j^r) \geq V_j^d - V_j^c,$$

which gives:

$$V_j^c \geq \frac{1}{K_j + 1}(V_j^d + K_j V_j^r).$$

On the left, this is the average return over stages of an always cooperating player, on the right this is the average return over stages of any deviating player. Therefore, for any $\{\pi_j^t\}_t \neq \{\pi_{\text{Folk}}^t\}_t$:

$$\mathbb{E}_{t \geq 0} [\mathcal{R}_j(\pi_{\text{Folk}}^t)] \geq \mathbb{E}_{t \geq 0} [\mathcal{R}_j(\pi_j^t, \pi_{\text{Folk}_{-j}}^t)].$$

■

Appendix B. Proof of Thm. 1

Proof Since σ is N-cyclic, any player i receives the same average return every N stage games:

$$\begin{aligned}
\mathbb{E}_{t \geq 0} [\mathcal{R}_i(\pi^t)] &= \frac{1}{N} \sum_{t=1}^N \mathcal{R}_i(\pi_{\sigma^t(i)}^\Sigma, \pi_{\sigma^t(-i)}^\Sigma) \\
&\quad (\text{the strategy is stationary}) \\
&= \frac{1}{N} \sum_{t=1}^N \mathcal{R}_{\sigma^t(i)}(\pi_i^\Sigma, \pi_{-i}^\Sigma) \\
&\quad (\text{the game is symmetric}) \\
&= \frac{1}{N} \sum_{t=1}^N \mathcal{R}_t(\pi_i^\Sigma, \pi_{-i}^\Sigma) \\
&\quad (\text{changing the order}).
\end{aligned}$$

Consequently, π_t maximizes the sum of returns at any t and the average return of the strategy is the same for all players. Now, imagine there is a strategy $\{\hat{\pi}^t\}_t$ such that:

$$\min_i \mathbb{E}_{t \geq 0} [\mathcal{R}_i(\hat{\pi}_i^t, \hat{\pi}_{-i}^t)] > \min_i \mathbb{E}_{t \geq 0} [\mathcal{R}_i(\pi_i^t, \pi_{-i}^t)].$$

In that case,

$$\begin{aligned}
\sum_i \mathbb{E}_{t \geq 0} [\mathcal{R}_i(\hat{\pi}_i^t, \hat{\pi}_{-i}^t)] &> N \min_i \mathbb{E}_{t \geq 0} [\mathcal{R}_i(\pi_i^t, \pi_{-i}^t)] \\
&= \sum_i \mathbb{E}_{t \geq 0} [\mathcal{R}_i(\pi_i^t, \pi_{-i}^t)],
\end{aligned}$$

which is in contradiction with the fact that π_t maximizes the sum of returns at any t . ■

Appendix C. Proof of Thm. 2

Proof Let first show that FCL converges to cooperation in self-play (when all agents are playing FCL). Given the fact that ϕ is deterministic, there are no defections. Let's focus on $\{\hat{\pi}^t\}_t = \{\pi^t | \phi(t) = \mathbf{True}\}$ the endless sub-process corresponding to exploration times. Let $\mathbb{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_N$. Clearly, for all $t, s \in \mathcal{S}$ and $a \in \mathbb{A}$, $\hat{\pi}^t(a|s) > 0$. Consequently, the convergence of Q -functions is given by the convergence of the classic Q -learning using $\pi(a^t = a | s^t = s) = \hat{\pi}^t(a|s)$ Melo (2001). Let Q^{c*} and Q_j^{r*} be the corresponding points of convergence. By construction:

- π^c , maximizing Q^{c*} , maximizes the sum of players returns.
- $\pi^* = (\pi_{\sigma^t(1)}^c \dots \pi_{\sigma^t(N)}^c)$ maximizes the min of players returns.
- $\pi^r(j)$, minimizing Q_j^{r*} , retaliates when player j deviates from π^* .

FCL decision rule at line 7 in Alg.1 corresponds to playing according to π^* when Q^c is close enough to Q^{c*} (when the difference between the max value and the second-max value is larger than twice an update size). Similarly, decision rule at line 11 corresponds to playing according to $\pi^r(j)$ when Q_j^r is close enough to Q_j^{r*} . Let T_Q be the smallest time after which both Q^c and Q_j^r are close enough to Q^{c*} and Q_j^{r*} so lines 7 and 11 correspond to playing according to π^* and $\pi^r(j)$. If explorations are stopped, FCL behaves as a cooperative equilibrium.

Now, let $A_j = \{\pi_j^t\}_t$ be a learning algorithm different than FCL and played by an agent j while all other players follow FCL ($A_{-j} = \{\pi_{-j}^t\}_t$). We will distinguish two situations:

- (a) $\forall s \in \mathcal{S}, \forall a_j \in \mathcal{A}_j, \forall a_{-j} \in \mathcal{A}_{-j}, \forall T, \exists t > T, \pi_j^t(a_j|s) \times \pi_{-j}^t(a_{-j}|s) > 0,$
- (b) $\exists s \in \mathcal{S}, \exists a_j \in \mathcal{A}_j, \forall a_{-j} \in \mathcal{A}_{-j}, \exists T, \forall t > T, \pi_j^t(a_j|s) \times \pi_{-j}^t(a_{-j}|s) = 0.$

In situation (a), all possible transitions are explored: the conditions for the convergence of FCL agents are met and they converge to a cooperative equilibrium, retaliating if the non-FCL learner is not cooperative. However, we must make sure that FCL's explorations are not spoiling the efficiency of retaliations: Let $\mathcal{R}_1 = \mathbb{E}_{t > T_\epsilon}[\mathcal{R}_j(A_j, A_{-j}^*)]$ be a deviating player's average return when FCL players are greedy (with probability $1 - \epsilon$), \mathcal{R}_2 be its average return when other agents are exploring (with probability ϵ , and \mathcal{R}_1^* and \mathcal{R}_2^* be the respective returns when no player deviates). Because the game is symmetric, we know that $\mathcal{R}_1 \leq \mathcal{R}_1^*$. If $\mathcal{R}_2 \leq \mathcal{R}_2^*$, then the average return of a deviating player is always smaller than if it does not deviate. Otherwise, we have $\mathcal{R}_2 + \mathcal{R}_1^* > \mathcal{R}_2^* + \mathcal{R}_1$ and by taking:

$$\epsilon = \frac{\mathcal{R}_1^* - \mathcal{R}_1}{\mathcal{R}_2 + \mathcal{R}_1^* - (\mathcal{R}_2^* + \mathcal{R}_1)}$$

we obtain, for all $t > T_\epsilon$:

$$(1 - \epsilon)\mathcal{R}_1 + \epsilon\mathcal{R}_2 \leq (1 - \epsilon)\mathcal{R}_1^* + \epsilon\mathcal{R}_2^*.$$

Consequently:

$$\mathbb{E}_{t > T_\epsilon}[\mathcal{R}_j(A_j(t), A_{-j}^*(t))] \leq \mathbb{E}_{t > T_\epsilon}[\mathcal{R}_j(A_j^*(t), A_{-j}^*(t))].$$

In situation (b), there is still a subset of state-action couples Ω_∞ that will be explored an infinite number of times. If all other players restrict their states and actions to the same subset (using $\pi_i(a|s) > 0 \Leftrightarrow (a, s) \in \Omega_\infty$) the induced sub-game is still symmetric and player j is exploring the whole sub-game an infinite number of times. Consequently, FCL can at least learn a cooperative equilibrium $\{\hat{\pi}_{\text{Folk}}^t\}_t$ based on a retaliation strategy $\pi_{\Omega_\infty}^{r,j}$ and a cooperative strategy $\pi_{\Omega_\infty}^*$ defined on Ω_∞ such that:

$$\forall \{\pi_j^t\}_t \neq \{\hat{\pi}_{\text{Folk}_t}^t\}, \mathbb{E}_{t \geq 0}[\mathcal{R}_j(\hat{\pi}_{\text{Folk}}^t)] = \mathbb{E}_{t \geq 0}[\mathcal{R}_j(\pi_{\Omega_\infty}^*)] \geq \mathbb{E}_{t \geq 0}[\mathcal{R}_j(\pi_j^t, \hat{\pi}_{\text{Folk}_{-j}}^t)].$$

Since $\pi_{\Omega_\infty}^*$ is necessarily sub-optimal to cooperate in the whole game, we have:

$$\forall \{\pi_j^t\}_t \neq \{\hat{\pi}_{\text{Folk}_t}^t\}, \mathbb{E}_{t \geq 0}[\mathcal{R}_j(\pi^*)] \geq \mathbb{E}_{t \geq 0}[\mathcal{R}_j(\pi_{\Omega_\infty}^*)] \geq \mathbb{E}_{t \geq 0}[\mathcal{R}_j(\pi_j^t, \hat{\pi}_{\text{Folk}_{-j}}^t)].$$

As a consequence, players can still retaliate and we can use the exact same argument than in (a) to obtain the desired statement. \blacksquare

References

- Itai Ashlagi, Dov Monderer, and Moshe Tennenholtz. Robust learning equilibrium. *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, 2006.
- Michael Bowling and Manuela Veloso. Rational and convergent learning in stochastic games. *Proceedings of the International joint conference on artificial intelligence*, 2001.
- Ronen I Brafman and Moshe Tennenholtz. Efficient learning equilibrium. *Advances in Neural Information Processing Systems*, 2003.
- Ronen I Brafman and Moshe Tennenholtz. Optimal efficient learning equilibrium: Imperfect monitoring in symmetric games. *Proceedings of the National Conference on Artificial Intelligence*, 2005.
- Caroline Claus and Craig Boutilier. The dynamics of reinforcement learning in cooperative multiagent systems. *Proceedings of the Association for the Advancement of Artificial Intelligence*, 1998.
- Partha Dasgupta, Eric Maskin, et al. The existence of equilibrium in discontinuous economic games. *Review of Economic Studies*, 1986.
- Jakob Foerster, Richard Y Chen, Maruan Al-Shedivat, Shimon Whiteson, Pieter Abbeel, and Igor Mordatch. Learning with opponent-learning awareness. *Proceedings of the International Conference on Autonomous Agents and MultiAgent Systems*, 2018a.
- Jakob N Foerster, Gregory Farquhar, Triantafyllos Afouras, Nantas Nardelli, and Shimon Whiteson. Counterfactual multi-agent policy gradients. In *Thirty-second AAAI conference on artificial intelligence*, 2018b.
- Edward Hughes, Joel Z Leibo, Matthew Phillips, Karl Tuyls, Edgar Dueñez-Guzman, Antonio García Castañeda, Iain Dunning, Tina Zhu, Kevin McKee, Raphael Koster, et al. Inequity aversion improves cooperation in intertemporal social dilemmas. *Advances in Neural Information Processing Systems*, 2018.
- Max Kleiman-Weiner, Mark K Ho, Joseph L Austerweil, Michael L Littman, and Joshua B Tenenbaum. Coordinate to cooperate or compete: abstract goals and joint intentions in social interaction. *Proceedings of Annual Conference of the Cognitive Science Society*, 2016.
- Joel Z Leibo, Vinicius Zambaldi, Marc Lanctot, Janusz Marecki, and Thore Graepel. Multi-agent reinforcement learning in sequential social dilemmas. *Proceedings of the Conference on Autonomous Agents and MultiAgent Systems*, 2017.
- Adam Lerer and Alexander Peysakhovich. Maintaining cooperation in complex social dilemmas using deep reinforcement learning. *arXiv preprint arXiv:1707.01068*, 2017.
- Alistair Letcher, Jakob Foerster, David Balduzzi, Tim Rocktäschel, and Shimon Whiteson. Stable opponent shaping in differentiable games. *Proceedings of the International Conference on Learning Representations*, 2018.

- Michael L Littman. Markov games as a framework for multi-agent reinforcement learning. *Proceedings of the International Conference on Machine Learning*, 1994.
- Michael L Littman. Friend-or-foe q-learning in general-sum games. *Proceeding of the International Conference on Machine Learning*, 2001.
- Michael L Littman and Peter Stone. A polynomial-time nash equilibrium algorithm for repeated games. *Decision Support Systems*, 2005.
- Francisco S Melo. Convergence of q-learning: A simple proof. *Institute Of Systems and Robotics, Tech. Rep*, 2001.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 2015.
- Enrique Munoz de Cote and Michael L. Littman. A polynomial-time Nash equilibrium algorithm for repeated stochastic games. *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, 2008.
- John Nash. Non-cooperative games. *Annals of mathematics*, 1951.
- Martin J Osborne and Ariel Rubinstein. *A course in game theory*. MIT press, 1994.
- Julien Pérolat, Joel Z Leibo, Vinicius Zambaldi, Charles Beattie, Karl Tuyls, and Thore Graepel. A multi-agent reinforcement learning model of common-pool resource appropriation. *Advances in Neural Information Processing Systems*, 2017.
- Alexander Peysakhovich and Adam Lerer. Consequentialist conditional cooperation in social dilemmas with imperfect information. *Proceedings of the International Conference on Learning Representations*, 2018.
- Lloyd S Shapley. Stochastic games. *Proceedings of the national academy of sciences*, 1953.
- Moshe Tennenholtz and Aviv Zohar. Learning equilibria in repeated congestion games. *Proceedings of The International Conference on Autonomous Agents and Multiagent Systems*, 2009.
- Steen Vester. *Symmetric Nash Equilibria*. PhD thesis, Master thesis from Ecole Normale Supérieure de Cachan, 2012.
- Christopher JCH Watkins and Peter Dayan. Q-learning. *Machine Learning*, 1992.
- Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 1992.