

On Fair Cost Sharing Games in Machine Learning

Ievgen Redko Charlotte Laclau

Univ Lyon, UJM-Saint-Etienne, CNRS, Institut d'Optique Graduate School

Laboratoire Hubert Curien UMR 5516

F-42023, Saint-Etienne, France

name.surname@univ-st-etienne.fr

Abstract

Machine learning and game theory are known to exhibit a very strong link as they mutually provide each other with solutions and models allowing to study and analyze the optimal behaviour of a set of agents. In this paper, we take a closer look at a special class of games, known as *fair cost sharing games*, from a machine learning perspective. We show that this particular kind of games, where agents can choose between selfish behaviour and cooperation with shared costs, has a natural link to several machine learning scenarios including collaborative learning with homogeneous and heterogeneous sources of data. We further demonstrate how the game-theoretical results bounding the ratio between the best Nash equilibrium (or its approximate counterpart) and the optimal solution of a given game can be used to provide the upper bound of the gain achievable by the collaborative learning expressed as the expected risk and the sample complexity for homogeneous and heterogeneous cases, respectively. We believe that the established link can spur many possible future implications for other learning scenarios as well, with privacy-aware learning being among the most noticeable examples.

Introduction

In recent years, machine learning community witnessed an increasing interest among the researchers towards the game theory, a sub-field of mathematics that studies the problem of decision-making in the presence of various types of constraints. This is no surprise as game theory proposes tools and a large variety of theoretical results allowing to determine optimal strategies of a set of agents that may have conflicting or collaborative objectives: a situation commonly encountered in many machine learning multi-objective optimization problems. This kind of problems, for instance, can be faced by game playing machine learning algorithms that seek to find an optimal trade-off between successfully passing the game level, earning the highest number of bonuses and doing all this in the fastest possible way. To this end, game theory has become a topic of ongoing interest in machine learning field that has already found its application in contributions related to numerous learning scenarios such as reinforcement learning (Peshkin et al. 2000; Hu and Wellman 2003; Claus and Boutilier 1998; Panait and Luke 2005), supervised learning (Freund and Schapire 1999; Shalev-Shwartz and Singer 2007a; 2007b;

Schuurmans and Zinkevich 2016), and adversarial classification (Liu and Chawla 2009; Brückner and Scheffer 2011; Dritsoula, Loiseau, and Musacchio 2017) to name a few.

In this paper, we take a closer look at *fair cost sharing games*, a special case of games that studies the optimal strategy of a set of agents when they can choose between a cooperative behaviour with a cost equally shared among them and a non-cooperative (also called “selfish”) behaviour. As an example of this game, one can consider a set of colleagues that face a choice between driving their cars to work and individually paying for it or taking the bus together and sharing the cost of the trip. We show that this particular game can be naturally related to collaborative learning, a problem often encountered in machine learning that consists in finding the best hypothesis for a set of data samples drawn from (possibly) different probability distributions and achieving a nearly optimal individual performance with respect to some task. In this setting, we associate each agent with a data sample that it can share with other agents in order to learn in a collaborative way on a larger concatenated sample or stick to what it has and learn on the sample available to it. We analyze this problem in several settings where the agents’ sample may or may not be drawn from the same probability distribution and where each agent may have a weight that corresponds to its potential benefit from collaboration and its contribution to it. For both cases considered, we propose a theoretical result that bounds the ratio between the overall cost of non-collaborative learning with respect to the collaborative one where the cost can be defined based on empirical risk achieved by the optimal hypothesis or on the sample complexity of the considered learning approach. To the best of our knowledge, this is the first contribution that establishes a connection between this class of games and collaborative learning and shows its usefulness in providing new theoretical guarantees for the latter.

The rest of this paper is organized as follows. We first introduce the related works that exist in the literature providing the theoretical analysis of the considered learning settings both in traditional and game-theoretical contexts. Then, we present necessary background definitions, on which we rely in the following sections, that introduce both basic and weighted versions of the fair cost sharing games and theoretical results established for them. We further proceed by first formally describing the considered setup and

then by presenting our main contributions based on it. The last section of this paper is devoted to conclusions and to the description of several future perspectives of this work.

Related Works

Despite a considerable amount of work situated at the intersection of game theory and machine learning mentioned in the previous section, few of them are related to the main contributions of this paper in terms of the quantities of interest that they analyze and the class of games that they consider. We present the related work structured with respect to these criteria below.

Mechanism design via machine learning Arguably, one of the first papers that analyzed a certain class of games using the concepts from statistical learning theory and sample complexity in particular was presented in (Balcan et al. 2005)¹. Their contribution considered a revenue-maximizing game where the goal is to find an optimal pricing function for a set of auction bidders, and consisted in showing that the optimal solution of this problem can be characterized using the techniques from statistical learning theory. We, however, consider a different class of games and, more importantly, undertake the opposite direction that consists in providing new results for collaborative machine learning problems using game-theoretical concepts. As this research direction is in general quite unrelated to ours, we refer the interested reader to (Liu, Chen, and Qin 2015, Related work) for a more complete up-to-date survey on the subject.

Statistical cost sharing Several recent papers (Balcan, Procaccia, and Zick 2015; Balkanski, Syed, and Vassilvitskii 2017) considered the class of cost sharing games that present the main subject of investigation of this paper. The main goal of (Balcan, Procaccia, and Zick 2015) was to define an algorithmic approach with strong theoretical guarantees that allows to calculate the cost-sharing function and define the optimal costs of agents based on it. The authors of (Balkanski, Syed, and Vassilvitskii 2017) improved the analysis provided in (Balcan, Procaccia, and Zick 2015) and also addressed the estimation of the Shapley value (Shapley 1953), a unique vector of cost shares that satisfies a set of natural axioms (Herzog, Shenker, and Estrin 1995). These papers are similar to ours as they relate the probably approximately correct (PAC) analysis (Valiant 1984) to cost sharing games. However, our work differs from these latter in two principal ways: (1) while cited papers aim to find an algorithmically optimal way to calculate the costs of collaboration for each agent satisfying several natural axioms, we consider a special case of cost sharing games with a fair division scheme; (2) similar with mechanism design contributions mentioned

¹This paper is a follow-up work of (Blum et al. 2003; Blum and Hartline 2005) where the same problem was considered in the context of online learning but without relying on the sample complexity results.

above, our paper aims at applying results from game theory to provide new insights for PAC analysis of collaborative learning, while (Balcan, Procaccia, and Zick 2015; Balkanski, Syed, and Vassilvitskii 2017) tailor traditional PAC analysis to study cost sharing games.

Strategyproof classification Strategy-proof classification problem studied in (Dekel, Fischer, and Procaccia 2010; Meir, Procaccia, and Rosenschein 2012) deals with a collaborative learning setup where a group of agents have a choice between reporting their true labels or falsifying them in order to achieve a better individual classifier. For this problem, the above-mentioned papers showed that, under some assumptions, the popular empirical risk minimizing (ERM) mechanism is truthful and optimal, *i.e.*, it encourages all agents to report their true labels and provides an approximately optimal solution for all of them. Our work is close to this line of research as it also considers empirical risk as a cost of collaboration for each agent and studies the general collaborative learning scenario in the PAC setting. Despite this similarity, the purpose of our work is different as we aim to study the gain possibly achievable by a collaborative learning algorithm: a question that was not addressed by either of these works.

Collaborative learning Finally, we briefly cover the contributions that establish theoretical guarantees for collaborative learning setting considered in this paper, where a set of agents aim at learning an accurate model simultaneously. The notion of collaborative learning is very vast and covers such areas as multi-task learning (Baxter 1997; Caruana 1997; Kumar and III 2012), multi-source domain adaptation (Ben-David et al. 2010; Mansour, Mohri, and Rostamizadeh 2009a; 2009b) and distributed learning (Balcan et al. 2012; Wang, Kolar, and Srerbo 2016). To this end, we note that our work is similar to that presented in (Blum et al. 2017) where the authors seek to establish the approximate value of the ratio between the sample complexity of non-collaborative learning and collaborative learning settings: a quantity denoted by them as overhead. The authors further propose an algorithm and a PAC analysis showing that its overhead is logarithmic. Our contribution completes this analysis with an upper-bound of the overhead and provides this result in a different, and arguably simpler, way. We elaborate more on the link between the two results in the section where the main contributions of our paper are presented.

Preliminary Knowledge

In this section, we present the preliminary knowledge related to the game-theoretic concepts that we use later in this paper. We start with a general description of a game and proceed by introducing a fair cost sharing game and some results obtained for it.

General Definitions

Given N , a set of K agents, S_i , a finite action space of agent $i \in N$ and c_i , a cost function of agent i , a game is defined

as a tuple

$$\mathcal{G} = \langle N, (\mathcal{S}_i), (c_i) \rangle.$$

We define the joint action space of the agents as $\mathcal{S} = \mathcal{S}_1 \times \dots \times \mathcal{S}_K$ and let the cost function c_i associated to an agent i be a mapping of a joint action $s \in \mathcal{S}$ to a real non-negative number, *i.e.*, $c_i : \mathcal{S} \rightarrow \mathbb{R}^+$. In this work, we assume that the social cost $c : \mathcal{S} \rightarrow \mathbb{R}^+$ is defined as the overall sum of agent's costs $\sum_{i=1}^K c_i$. The optimal social cost for this scenario is given by the strategy minimizing the cost c :

$$\text{OPT}(\mathcal{G}) = \min_{s \in \mathcal{S}} c(s).$$

In this game, we say that a joint action $s \in \mathcal{S}$ (also called strategy) is a *pure Nash Equilibrium* (Nash 1950; 1951) if no agent $i \in N$ can benefit from unilaterally deviating to another action. Denoting by $N(\mathcal{G})$ the set of Nash equilibria of the game \mathcal{G} , we further define two key quantities related to games as follows.

Definition 1. *The Price of Anarchy (PoA) is the ratio of the worst Nash equilibrium to the social optimum. It measures how the efficiency of a system deteriorates due to a selfish behavior of the agents of the game, *i.e.*,*

$$\text{PoA}(\mathcal{G}) = \max_{s \in N(\mathcal{G})} c(s) / \text{OPT}(\mathcal{G}).$$

Definition 2. *The Price of Stability (PoS) is the ratio of the best Nash equilibrium to the social optimum, *i.e.*,*

$$\text{PoS}(\mathcal{G}) = \min_{s \in N(\mathcal{G})} c(s) / \text{OPT}(\mathcal{G}).$$

The motivation behind introducing PoS in addition to PoA stems from the fact that this latter can be very large, making it uninformative in practice. As we show it below, this is the case for a particular class of games studied in this paper.

We also note that some classes of games do not admit a pure Nash equilibrium (Anshelevich et al. 2003) so that an approximate Nash equilibrium should be considered. This latter can be defined as a strategy for which no agent can decrease its cost by more than an α multiplicative factor from unilaterally deviating to another action. In this case, the quantities PoS and PoA are defined with respect to the α -approximate Nash equilibrium in the same manner.

Fair Cost Sharing Game

In this paper, we focus on a specific class of games, referred to as *fair cost sharing games*. Such a game take place in a graph $G = (V, E)$ with a set of K agents, where each edge $e \in E$ carries a non-negative cost γ_e , and each agent i has source node $s_i \in V$ and destination node $t_i \in V$ that it tries to connect. We denote by \mathcal{S}_i the set of paths taken by agent i in order to connect s_i to t_i . Outcomes of the game correspond to path vectors $s = (P_1, \dots, P_K)$, with each agents choosing a single path $P_i \in \mathcal{S}_i$. We further denote by x_e the number of agents whose strategy contains edge e .

One can think of γ_e as the fixed cost of building the edge e , and this cost is independent of the number of agents that use the edge. Therefore, if more than one agent use an edge e in their chosen paths, *i.e.*, $x_e > 1$, then they share the edge's fixed cost γ_e . In a *fair cost sharing game*, we assume that

the cost is split equally among the agents meaning that the cost to any agent i is

$$c_i(s) = \sum_{e \in P_i} \frac{\gamma_e}{x_e}.$$

Then, the objective is to minimize the total cost of the formed network defined as

$$c(P_1, \dots, P_K) = \sum_{e \in \bigcup_i P_i} \gamma_e.$$

As an illustrative example, one can consider a special instance of a fair cost sharing game, referred to as *opting out*, presented in Figure 1. Here, the K agents have distinct sources s_1, \dots, s_K , but a common destination t . In order to reach t , they have two options: (1) meeting at a rendezvous point v and continuing together to t , resulting in a joint cost of $1 + \varepsilon$, for some small $\varepsilon > 0$ or (2) taking the direct $s_i - t_i$ path individually. In this case, each agent i incurs a cost of $\frac{1}{i}$ for its opt-out strategy leading to a unique Nash equilibrium with cost $\mathcal{H}_K = \sum_{i=1}^K \frac{1}{i}$. However, one can clearly observe that the optimal solution in this game would be for all players to travel through the rendezvous point for an overall total cost of $1 + \varepsilon$ incurring an individual cost of $\frac{1+\varepsilon}{K}$ for each agent. Furthermore, in the worst case, Nash equilibria of this game can be very expensive, so that the PoA becomes as large as K . To see this, we can consider a graph with common source and destination nodes for all agents and two parallel edges of cost 1 and K between them. In this case, the worst equilibrium corresponds to all players choosing the more expensive edge and paying K times the cost of the optimal solution. In its turn, PoS can be bounded due to the following theorem.

Theorem 1 ((Anshelevich et al. 2004)). *The PoS for pure Nash equilibria in fair cost sharing games is at most $\mathcal{H}_K = \Theta(\log K)$, where $\mathcal{H}_K = 1 + \frac{1}{2} + \dots + \frac{1}{K}$ and this bound is tight.*

This theorem provides us with a trade-off that can exist between the cost of selfish behaviour related to a pure Nash equilibrium for a set of agents and that of an optimal social cost. As shown in (Anshelevich et al. 2004), this bound is not vacuous as one may always find an example of a game where the ratio between the two is exactly $\Theta(\log K)$.

Note that the fair cost sharing game presented above admits that every agent pays the same cost for using the shared edge, even though in many real-world applications one may expect the cost to be shared among agents depending on their weights $\{w_i\}_{i=1}^K$. This weight, for instance, can be related to the contribution of a given agent in the collaboration. This scenario leads to a *weighted fair cost sharing game*, where the cost of each agent i is proportional to its weight w_i and can be calculated as follows:

$$c_i(s) = \sum_{e \in P_i} \gamma_e \frac{w_i}{W_e}$$

with W_e denoting the total weight of the players that select a path containing e .

For this particular case, the pure Nash equilibrium exists only in games with 2-players (Anshelevich et al. 2004)




Figure 1: Example of cost-sharing network known as opting out. Here the cost of Nash equilibrium is $\mathcal{H}_K = \sum_{i=1}^K \frac{1}{i}$ while the optimal cost is $1 + \epsilon$.

and thus α -approximate Nash equilibria are usually considered. Before presenting the theorem that bounds the PoS of weighted fair cost sharing games, we first assume that for all i , $w_i \geq 1$ and denote by $w_{\max} = \max_{i \in [1, \dots, K]} w_i$ the maximum

weight across all agents. We further let $W = \sum_{i=1}^K w_i$ be the overall sum of weights of all agents. The desired result can be now stated as follows.

Theorem 2 ((Chen and Roughgarden 2009)). *For $\alpha = \Omega(\log w_{\max})$, every weighted fair cost sharing game admits a $\mathcal{O}(\alpha)$ -approximate Nash equilibrium for which the PoS is at most $\mathcal{O}\left(\frac{\log(W)}{\alpha}\right)$.*

With these two theorems, we are now ready to present our main contributions.

Main Contributions

In this section, we present our main contribution that consists in showing how different learning paradigms can be seen as instances of fair cost sharing games. Our goal here is two-fold and consists in showing that: (1) representing different learning scenarios as instances of fair cost sharing games establishes a connection between the statistical learning theory and the game theory; (2) this connection can be used to inherit some important guarantees established in the rich literature on game theory. With this in mind, we now proceed to a formal description of the considered setup.

Problem Setup

Let us consider a set of K agents $\{a_i\}_{i=1}^K$, where each agent has access to a learning sample $S^i = \{(\mathbf{x}_j^{(i)}, y_j^{(i)})\}_{j=1}^{m_i}$ of size m_i , for all $i \in [1, \dots, K]$. For each i , we assume that S^i is drawn i.i.d. from a probability distribution \mathcal{D}^i defined over a product space $\mathbf{X} \times Y$, where $\mathbf{X} \subseteq \mathbb{R}^d$ and Y is an output space that can be equal to $\{0, 1\}$ in case of binary classification. In practice, S^i can be given by a collection of images, while classes $\{0, 1\}$ may define the presence or absence of a certain object on an image. For a convex loss

function $\ell : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}_+$ and a sample S^i , we define the true and empirical risks for each $i \in [1, \dots, K]$ as follows:

$$\begin{aligned} \mathbb{R}_{\mathcal{D}^i}(h(\mathbf{x}), y) &= \mathbf{E}_{(\mathbf{x}, y) \sim \mathcal{D}^i} [\ell(h(\mathbf{x}), y)], \\ \mathbb{R}_{\hat{\mathcal{D}}^i}(h(\mathbf{x}), y) &= \mathbf{E}_{(\mathbf{x}, y) \sim \hat{\mathcal{D}}^i} [\ell(h(\mathbf{x}), y)] \\ &= \frac{1}{m_i} \sum_{j=1}^{m_i} \ell(h(\mathbf{x}_j^{(i)}), y_j^{(i)}), \end{aligned}$$

where $\hat{\mathcal{D}}^i = \frac{1}{m_i} \sum_{j=1}^{m_i} \delta_{\mathbf{x}_j^{(i)}}$ is an empirical distribution associated with \mathcal{D}^i and $h \in \mathcal{H}$ is a hypothesis from some hypothesis space \mathcal{H} such that $h : \mathbf{X} \rightarrow Y$. As an example, one may consider \mathcal{H} as a space of linear functions so that h would be a hyperplane that separates the two classes. We now define the risk-minimizing hypothesis as follows: let us denote by $h_{S^i}^* = \operatorname{argmin}_{h \in \mathcal{H}} \mathbb{R}_{\hat{\mathcal{D}}^i}(h(\mathbf{x}), y)$ and $h_{\mathcal{D}^i}^* = \operatorname{argmin}_{h \in \mathcal{H}} \mathbb{R}_{\mathcal{D}^i}(h(\mathbf{x}), y)$ for all $i \in [1, \dots, K]$ the empirical and true risk-minimizing hypotheses, respectively. We further denote by $\hat{\mathbb{R}}_{S^i}^*(\mathcal{H}) = \mathbb{R}_{\hat{\mathcal{D}}^i}(h_{S^i}^*(\mathbf{x}), y)$ and $\mathbb{R}_{\mathcal{D}^i}^*(\mathcal{H}) = \mathbb{R}_{\mathcal{D}^i}(h_{\mathcal{D}^i}^*(\mathbf{x}), y)$ the ideal empirical and true risks, respectively.

Collaborative Learning with Homogeneous Sources

In order to present our first result, we start by considering a traditional collaborative learning setting where agents have access to learning samples S^i drawn from the same underlying probability distribution $\mathcal{D} = \mathcal{D}^i, \forall i \in [1, \dots, K]$. For instance, this scenario can occur in practice when several hospitals build predictive models based on their collected data consisting of annotated MRI scans: if the scanners used to produce images are the same, we can suppose that the statistical distribution of images related to the same organ is also highly similar. Bearing in mind the high cost of manual labeling of MRI scans required to increase the sample size so that a low-error hypothesis can be learned, the hospitals may think of joining their forces and pooling their labeled samples together. In this case, the goal of our analysis would be to derive a bound on the ratio between the overall performance achieved by individual agents that learn on a limited sample available and the performance of a hypothesis obtained using a larger sample $S = \bigcup_{i=1}^K S^i$.

In order to make the considered problem more realistic, we attribute weights to each agent reflecting the number of labeled instances that it provides to the collaborative learning algorithm. Intuitively, the cost of collaboration for agents that have large data samples should be smaller as they benefit less from collaboration due to their capacity of being able to learn a good classifier on their own. To this end, we define the weights w_i of agents for all $i \in [1, \dots, K]$ as a ratio $\frac{m_{\max}}{m_i}$, where $m_{\max} = \max_{i \in [1, \dots, K]} m_i$. This definition ensures that the weight of the agent having access to the largest sample is equal to 1, while for all the others it is greater or equal than 1. We can now state the following theorem.

Theorem 3. *Assume that for all $i \in [1, \dots, K]$, $\mathcal{D} = \mathcal{D}^i$. Let $h_S^* = \operatorname{argmin}_{h \in \mathcal{H}} \mathbb{R}_{\hat{\mathcal{D}}}(h(\mathbf{x}), y)$, where $\hat{\mathcal{D}}$ is an empirical*

distribution associated with the sample $S = \bigcup_{i=1}^K S^i$, such that $|S| = \sum_{i=1}^K m_i = m$ for a hypothesis space \mathcal{H} . Let $\hat{\mathbb{R}}_{\mathcal{D}_S^i}^*(\mathcal{H}) = \mathbf{E}_{(\mathbf{x}, y) \sim \hat{\mathcal{D}}^i} [\ell(h_S^*(\mathbf{x}), y)]$ and assume further that $\sum_{i=1}^K \hat{\mathbb{R}}_{S^i}^*(\mathcal{H}) \geq \sum_{i=1}^K \hat{\mathbb{R}}_{\mathcal{D}_S^i}^*(\mathcal{H}) \geq \frac{1}{\alpha} \max_{i \in [1, \dots, K]} \hat{\mathbb{R}}_{S^i}^*(\mathcal{H})$ for some $\alpha \geq 0$ with $\hat{\mathbb{R}}_{S^i}^*(\mathcal{H}) > 0$ for all $i \in [1, \dots, K]$. Then, the following holds:

$$\frac{\sum_{i=1}^K \hat{\mathbb{R}}_{S^i}^*(\mathcal{H})}{\sum_{i=1}^K \hat{\mathbb{R}}_{\mathcal{D}_S^i}^*(\mathcal{H})} \leq \mathcal{O} \left(\frac{\log \left(\sum_{i=1}^K \frac{m_{\max}}{m_i} \right)}{\alpha} \right).$$

Proof. The main idea of our proof is to show that this particular learning setting can be represented as an instance of a weighted fair cost sharing game. Once this is done, we can simply apply Theorem 2 to obtain the desired result.

To this end, we start by considering the construction represented in Figure 2 and proceed by defining the nodes and edges with their associated costs as follows. First, we let the nodes a_i correspond to source nodes of agents $\{a_i\}_{i=1}^K$ with their respective learning samples S^i . The node \mathbf{L} corresponds to the destination node where a risk-minimizing classifier is learned, while the node \mathbf{P} corresponds to pooling the data from the incoming edges. In this game, each agent a_i has a choice between learning a classifier using its own available sample S^i by taking the edge $a_i - \mathbf{L}$ or pooling it with other agents by choosing the path $a_i - \mathbf{P} - \mathbf{L}$. This latter choice stands for learning a classifier that minimizes the overall error on the union of their samples with an individual cost of $c^* \frac{w_i}{\sum_{i=1}^K w_i}$. In this case, we can define the costs of edges as follows: we assume that the cost of taking the edge from node a_i to \mathbf{L} has a cost of the optimal risk achievable by minimizing it over the observable sample S^i . Thus, we write for all $i \in [1, \dots, K]$, $c_i = \hat{\mathbb{R}}_{S^i}^*(\mathcal{H})$.

As we consider a scenario where all the data distributions \mathcal{D}^i are the same, it is reasonable to further assume that the price of pooling the data is equal to 0 for each agent as it does not require reducing the discrepancy between the different agents' distributions. In this case, the price of taking the edge between each node a_i and \mathbf{P} is set to 0.

The price of learning in a collaborative way can be characterized by the sum of optimal risks achieved for each agent with respect to the hypothesis minimizing the risk on sample S . Consequently, we define it by letting $c^* = \sum_{i=1}^K \hat{\mathbb{R}}_{\mathcal{D}_S^i}^*(\mathcal{H})$. As in the classical fair cost sharing game, each agent has a preference for choosing a ‘‘selfish’’ non-collaborative strategy that consists in learning using its own sample when the inequality $\sum_{i=1}^K \hat{\mathbb{R}}_{\mathcal{D}_S^i}^*(\mathcal{H}) \geq \frac{1}{\alpha} \max_{i \in [1, \dots, K]} \hat{\mathbb{R}}_{S^i}^*(\mathcal{H})$ holds. This latter condition corresponds to a α -approximate Nash equilibrium of this game that has a cost of $\sum_{i=1}^K \hat{\mathbb{R}}_{\mathcal{D}_S^i}^*(\mathcal{H})$. From the assumption of the theorem, the optimal solution, however, is to learn on a bigger sample by following the path $a_i - \mathbf{P} - \mathbf{L}$ with a cost of $\sum_{i=1}^K \hat{\mathbb{R}}_{\mathcal{D}_S^i}^*(\mathcal{H})$ that is shared by all agents. Applying Theorem 2, we can bound the ratio between the overall cost




Figure 2: **(Theorem 3)** Collaborative learning as a fair cost sharing game with multiple agents corresponding to different data sources generated from the same probability distributions. Here, for all $i \in [1, \dots, K]$, $c_i = \hat{\mathbb{R}}_{S^i}^*(\mathcal{H})$ while $c^* = \sum_{i=1}^K \hat{\mathbb{R}}_{\mathcal{D}_S^i}^*(\mathcal{H})$; **(Theorem 4)** Collaborative learning with data sources generated from different probability distributions. Here, for all $i \in [1, \dots, K]$, $c_i = m_{\epsilon, \delta}^i = m_{\epsilon, \delta}$ while $c^* = m_{\epsilon, \delta}^*$.

achieved at α -approximate Nash equilibrium and the optimal solution as follows:

$$\frac{\sum_{i=1}^K \hat{\mathbb{R}}_{\mathcal{D}_S^i}^*(\mathcal{H})}{\sum_{i=1}^K \hat{\mathbb{R}}_{S^i}^*(\mathcal{H})} \leq \mathcal{O} \left(\frac{\log \left(\sum_{i=1}^K \frac{m_{\max}}{m_i} \right)}{\alpha} \right).$$

□

The result established in this theorem states that for a considered learning scenario where additional data coming from K different sources can help to learn a better classifier, the gap between the sum of individual performances of all agents and that achieved by the collaborative learning approach increases when there exist important differences between the sample sizes of different agents. This implication is not trivial but rather intuitive as we may expect to obtain an improved performance at least for those agents who possess little data due to the data provided by other agents with bigger data samples. Note also that the statement of the theorem can be further simplified if $\alpha = 1$, *i.e.* α -approximate Nash equilibrium coincides with the pure one, but this latter exists only when $K = 2$.

Before proceeding to the analysis of a more general collaborative learning scenario, we first briefly comment on the main assumption of the theorem given by the inequality

$$\sum_{i=1}^K \hat{\mathbb{R}}_{S^i}^*(\mathcal{H}) \geq \sum_{i=1}^K \hat{\mathbb{R}}_{\mathcal{D}_S^i}^*(\mathcal{H}) \geq \frac{1}{\alpha} \max_{i \in [1, \dots, K]} \hat{\mathbb{R}}_{S^i}^*(\mathcal{H}). \quad (1)$$

First, we note that the left-hand side of (1) implies that a hypothesis learned on a bigger sample performs better on each individual sample of all agents a_i . Even though it

can be violated in some real-world scenarios, this assumption is rather intuitive in the framework of supervised learning where the generalization capacity tends to increase with the increasing sample size for a fixed hypothesis space. As for the right-hand side, the inequality $\sum_{i=1}^K \hat{\mathbb{R}}_{\mathcal{D}_i^S}^*(\mathcal{H}) \geq \frac{1}{\alpha} \max_{i \in [1, \dots, K]} \hat{\mathbb{R}}_{S^i}^*(\mathcal{H})$ means that the overall risk related to collaborative learning is higher than the worst performance over all agents multiplied by a factor of $\frac{1}{\alpha}$. It further implies that for all $i \in [1, \dots, K]$, $c_i \leq \alpha c^*$ meaning that the dominant strategy for all agents corresponding to the α -approximate Nash equilibrium is to learn using their own sample. This assumption restricts agents from deviating from their non-collaborative strategy but, as it was seen from the example presented in Figure 1, does not imply that the individual costs of all agents c_i are lower than the cost of a sharing edge c^*/K . In order to see this, one can consider an arbitrary vector $[c_1, \dots, c_K]$ with $\forall i \in [1, \dots, K], c_i > 0$ and denote by c_{\max} its maximum element. Given the condition $\frac{1}{\alpha} c_{\max} \leq \sum_{i=1}^K c_i^*$, one can always find $c^* = (c_1^*, \dots, c_K^*)$ such that $\forall i \in [1, \dots, K], c_i > c_i^*$. Indeed, we can set $c_1^* = \frac{1}{\alpha} (c_{\max} - \epsilon_1)$ for some $\epsilon_1 > 0$ and let $c_i^* = \frac{1}{\alpha} \left(\frac{\epsilon_1 + \epsilon_2}{K-1} \right)$, $\forall i \in [2, \dots, K]$, where the value of $\epsilon_2 > 0$ can be made infinitely small. In this case, the condition $\frac{1}{\alpha} c_{\max} \leq \sum_{i=1}^K c_i^*$ is verified and the values ϵ_1 and ϵ_2 can be chosen to ensure that $\forall i \in [1, \dots, K], c_i > c_i^*$ giving the desired result.

Finally, we note that the inequality established in this theorem holds even if agents are assumed to be self-interested and may lie about their true labels. This follows from the results established in (Dekel, Fischer, and Procaccia 2010) where the authors prove that costs defined as an empirical risk of the best hypothesis consistent with the learning sample encourage the agents to tell the truth about their labels. This, however, is only the case when \mathcal{H} is a space of constant or homogeneous linear functions over \mathbb{R}^d .

Collaborative Learning with Heterogeneous Sources

For our first result, we considered a setting where data distributions that generated individual samples of all agents are the same. This assumption, however, is often violated in practice as different data sources may provide data samples with important statistical differences. Consider the previous example with hospitals and assume now that the MRI scans are acquired on scanners with varying resolutions and sizes of the resulting images, thus leading to a shift in the statistical distribution between the different acquired samples. In this case, we fall into the category of the so-called transfer learning algorithms that may take the form of multi-task learning, domain adaptation or distributed learning problems. In what follows, we refer to any instance of a learning problem with heterogeneous samples following different probability distributions as to a general collaborative learning problem. In this scenario, a quantity of interest that one may want to quantify is the ratio between the number of samples that are needed to learn a good hypothesis for each agent and that required by a given collaborative algorithm to

produce a hypothesis (or a set of hypotheses) that performs well on all of them. In the context of PAC learnability, this can be formalized as follows.

Definition 3. Let \mathcal{H} be a hypothesis class of VC dimension d . We say that a hypothesis $h \in \mathcal{H}$ allows to (ϵ, δ) -solve a learning problem $(\mathcal{H}, \mathcal{D})$ if for any $\epsilon, \delta > 0$ with probability $1 - \delta$ the following holds:

$$\Pr_{\mathbf{x} \sim \mathcal{D}} [\mathbb{R}_{\mathcal{D}^i}(h(\mathbf{x}), y) \leq \epsilon] \geq 1 - \delta.$$

Let us now define a general collaborative learning problem as a 2-tuple $(\mathcal{H}, \{\mathcal{D}^i\}_{i=1}^K)$. We say that a hypothesis $h \in \mathcal{H}$ allows to (ϵ, δ) -solve a learning problem $(\mathcal{H}, \{\mathcal{D}^i\}_{i=1}^K)$ if with probability $1 - \delta$, $\mathbb{R}_{\mathcal{D}^i}(h(\mathbf{x}), y) \leq \epsilon$ for all $i \in [1, \dots, K]$. We further define the sample complexity $m_{\epsilon, \delta}^i$ as the size of sample S^i drawn from \mathcal{D}^i required by h to (ϵ, δ) -solve a problem $(\mathcal{H}, \mathcal{D}^i)$. We assume that \mathcal{H} is fixed for all individual agents so that, as shown in (Anthony and Bartlett 2009), the sample complexity becomes equal to $m_{\epsilon, \delta}^i = m_{\epsilon, \delta} = \mathcal{O}\left(\frac{1}{\epsilon} (d \log\left(\frac{1}{\epsilon} + \frac{1}{\delta}\right))\right)$, $\forall i \in [1, \dots, K]$. As an example of a collaborative learning algorithm, we consider (Blum et al. 2017, Algorithm 2)² and denote it by \mathcal{L} in what follows.

In this setting, we can now state the following theorem.

Theorem 4. Let \mathcal{H} be a hypothesis space of VC dimension d . Let $m_{\epsilon, \delta}^*$ be a sample complexity required to (ϵ, δ) -solve the collaborative learning problem $(\mathcal{H}, \{\mathcal{D}^i\}_{i=1}^K)$ by a hypothesis $h \in \mathcal{H}$ outputted by \mathcal{L} . Then, the following holds:

$$\frac{m_{\epsilon, \delta}}{m_{\epsilon, \delta}^*} \leq \frac{\Theta(\log(K))}{K}.$$

Proof. Similar to the previous theorem, the idea behind our proof is to represent a collaborative learning problem $(\mathcal{H}, \{\mathcal{D}^i\}_{i=1}^K)$ as a fair cost sharing game given in Figure 2. To this end, we let the nodes a_i correspond to the source nodes of agents $\{a_i\}_{i=1}^K$ with their respective distributions \mathcal{D}^i and let the pooling node be the same as before. The destination node \mathbf{L} is the state where every problem $(\mathcal{H}, \{\mathcal{D}^i\})$ is (ϵ, δ) -solved. The edge between node \mathbf{P} and \mathbf{L} corresponds to applying \mathcal{L} on the received samples from node \mathbf{P} . Each agent has a choice between generating a sample of size $m_{\epsilon, \delta}$ to (ϵ, δ) -solve their problem or generating a sample for a collaborative learner \mathcal{L} . In the proposed setting, we can set the costs of individual edges $c_i = m_{\epsilon, \delta}^i = m_{\epsilon, \delta}$, $\forall i \in [1, \dots, K]$. The shared edge thus has a cost of $c^* = m_{\epsilon, \delta}^*$, where the $m_{\epsilon, \delta}^*$ examples are drawn from a uniform mixture distribution $\frac{1}{K} \sum_{i=1}^K \mathcal{D}^i$. In order to learn a hypothesis in a collaborative setting, we consider the algorithm presented in (Blum et al. 2017, Algorithm 2) that has a property of (ϵ, δ) -solving a collaborative problem $(\mathcal{H}, \{\mathcal{D}^i\}_{i=1}^K)$ with $\frac{m_{\epsilon, \delta}^*}{m_{\epsilon, \delta}} = \mathcal{O}(\log^2(K))$ when $K = \mathcal{O}(d)$. In this case, agents $\{a_i\}_{i=1}^K$ tend to prefer paying a cost of

²For the sake of completeness, we provide the description of this algorithm and the theoretical result related to it in the Supplementary material.

$m_{\epsilon,\delta} \leq m_{\epsilon,\delta}^*$ for (ϵ, δ) -solving $(\mathcal{H}, \{\mathcal{D}^i\})$, $\forall i \in [1, \dots, K]$ thus leading to the Nash equilibrium with a cost equal to $Km_{\epsilon,\delta}$. However, the optimal solution is to learn in a collaborative way with a cost $m_{\epsilon,\delta}^* < Km_{\epsilon,\delta}$. As before, we can now bound the ratio between the Nash equilibrium and the optimal cost using Theorem 1 yielding the final result:

$$\frac{m_{\epsilon,\delta}}{m_{\epsilon,\delta}^*} \leq \frac{\Theta(\log(K))}{K}.$$

□

Remark 1. *Note that the established result can be proved in a more general setting without specifying the collaborative algorithm used to output a hypothesis. In this case, one would need to make an assumption similar to (1) that would restrict agents from choosing the shared path and ensure that the overall sample complexity of collaborative learning is smaller than the sum of sample complexities of each agent. This assumption, however, is almost always verified in practice as one can barely hope to find a collaborative algorithm that can (ϵ, δ) -solve K problems with a sample complexity smaller than that of each individual problem. The optimality condition, $Km_{\epsilon,\delta} > m_{\epsilon,\delta}^*$, is also far from being restrictive as it means that the chosen collaborative algorithm should perform at least a little bit better than a naive algorithm that takes a number of samples from each agent large enough to (ϵ, δ) -solve each problem $(\mathcal{H}, \{\mathcal{D}^i\})$.*

Remark 2. *Contrary to the first result obtained for homogeneous data sources, here we consider a simple fair cost sharing game with equal weights of all agents. As sample complexity remains the same for a fixed hypothesis space, it is reasonable to assume that agents participate equally in sharing the cost of the common edge. In the previous case, however, the sum of empirical errors defining the cost c^* consisted of terms that were not necessarily equal and thus the individual cost of each agent taking the shared path was allowed to be unequal too.*

It is import to underline that this result establishes a lower bound for the ratio between the collaborative and individual sample complexities (referred to as overhead) considered in (Blum et al. 2017). As mentioned in the proof, (Blum et al. 2017, Theorem 3.1) states that an algorithm for collaborative learning proposed in their paper has a sample complexity that is only $\mathcal{O}(\log^2(K))$ larger than the individual sample complexity of solving each problem $(\mathcal{H}, \{\mathcal{D}^i\})$. Obviously, in this case it achieves an exponential improvement over a naive algorithm that takes a number of samples from each agent large enough to (ϵ, δ) -solve each problem $(\mathcal{H}, \{\mathcal{D}^i\})$. The two results can be further combined in a unified statement that takes the following form:

$$\frac{K}{\Theta(\log(K))} \leq \mathcal{O}(\log^2(K)) = \frac{m_{\epsilon,\delta}^*}{m_{\epsilon,\delta}}.$$

Both this unification and the previous theorem highlight the practical interest in expressing various learning problems as instances of fair cost sharing game as it allows to provide an analysis of learning settings in a completely different, and arguably simpler, way.

Conclusions and Future Perspectives

In this paper, we considered collaborative learning problem widely presented in machine learning as an instance of fair cost sharing games. In order to relate the two, we showed how a collaborative learning scenario can be represented as a graph of a fair cost sharing game. The established link allowed us to analyze collaborative learning problem in two settings where the learning samples available to agents were assumed to be drawn from the same or different probability distributions. For these two cases, we proved two distinct theoretical results where the first one bounds the ratio between the optimal empirical risk of non-collaborative and collaborative learning approaches, while the second establishes the same kind of bound for their respective sample complexities. The proposed analysis naturally completes a previous work on the subject and brings a new point of view for the general collaborative learning scenario. To the best of our knowledge, there were no other attempts in analyzing this particular learning setting using the results established for fair cost sharing games in game theory.

Due to the rich body of literature on fair cost sharing games that exist in algorithmic game theory, the possible future perspectives of this work are many. Among them, one of the most important ones is to consider a fair cost sharing game where pooling the data has a certain cost for all agents that is different from 0. Indeed, in many real-world situations, agents represented by industrial actors may want to protect their data from its disclosure either to the learner or to the other agents. This conflict of interests is often modeled as a privacy-aware learning paradigm where the optimized objective function considers two terms representing the trade-off between the privacy of the data used for learning and the optimal learning itself. Fair cost sharing games offer a very intuitive way of modeling the above-mentioned situation by the means of imposing a non-zero cost on the edges that allow the agents to pool data together. In this case, the analysis proposed in this paper may be extended to study the possible gain of collaborative learning with privacy constraints and can be done, for instance, by using a traditional framework of differential privacy (Dwork 2006) and PAC learnability results related to it (Shiva Prasad Kasiviswanathan and Smith 2011; Dwork and Roth 2014). On the other hand, this setup can be also potentially applied to model the multi-source domain adaptation problem where the source nodes correspond to the available source domains, while the destination node is given by the target domain. In this case, the direct link between source nodes and the target one may correspond to the error achieved by a direct application of a source classifier to the target data usually considered as a weak baseline in domain adaptation. On the contrary, the shared edge would correspond to minimizing the combined weighted error over all sources in the spirit of (Blitzer et al. 2007) so that the cost would be defined based on the generalization bounds established by the authors of this paper. Both these research lines merit a thorough future investigation as they aim to tackle two important drawbacks of learning phenomenon such as its lack of generalizing capacity across different domains and the rising concerns regarding the privacy preservation.

References

- Anshelevich, E.; Dasgupta, A.; Tardos, E.; and Wexler, T. 2003. Near-optimal network design with selfish agents. In *STOC*, 511–520.
- Anshelevich, E.; Dasgupta, A.; Kleinberg, J.; Tardos, E.; Wexler, T.; and Roughgarden, T. 2004. The price of stability for network design with fair cost allocation. In *FOCS*, 295–304.
- Anthony, M., and Bartlett, P. L. 2009. *Neural Network Learning: Theoretical Foundations*. New York, NY, USA: Cambridge University Press, 1st edition.
- Balcan, M. F.; Blum, A.; Hartline, J. D.; and Mansour, Y. 2005. Mechanism design via machine learning. In *FOCS*, 605–614.
- Balcan, M.; Blum, A.; Fine, S.; and Mansour, Y. 2012. Distributed learning, communication complexity and privacy. In *COLT*, 26.1–26.22.
- Balcan, M.-F.; Procaccia, A. D.; and Zick, Y. 2015. Learning cooperative games. In *IJCAI*, 475–481.
- Balkanski, E.; Syed, U.; and Vassilvitskii, S. 2017. Statistical cost sharing. In *NIPS*, 6221–6230.
- Baxter, J. 1997. A bayesian/information theoretic model of learning to learn via multiple task sampling. *Machine Learning* 28(1):7–39.
- Ben-David, S.; Blitzer, J.; Crammer, K.; Kulesza, A.; Pereira, F.; and Vaughan, J. W. 2010. A theory of learning from different domains. *Machine Learning* 79(1-2):151–175.
- Blitzer, J.; Crammer, K.; Kulesza, A.; Pereira, F.; and Wortman, J. 2007. Learning bounds for domain adaptation. In *NIPS*, 129–136.
- Blum, A., and Hartline, J. D. 2005. Near-optimal online auctions. In *SODA*, 1156–1163.
- Blum, A.; Kumar, V.; Rudra, A.; and Wu, F. 2003. Online learning in online auctions. In *SODA*, 202–204.
- Blum, A.; Haghtalab, N.; Procaccia, A. D.; and Qiao, M. 2017. Collaborative pac learning. In *NIPS*, 2392–2401.
- Brückner, M., and Scheffer, T. 2011. Stackelberg games for adversarial prediction problems. In *ACM SIGKDD*, 547–555.
- Caruana, R. 1997. Multitask learning. *Machine Learning* 28(1):41–75.
- Chen, H.-L., and Roughgarden, T. 2009. Network design with weighted players. *Theoretical Computer Science* 45(2):302–324.
- Claus, C., and Boutilier, C. 1998. The dynamics of reinforcement learning in cooperative multiagent systems. In *AAAI*, 746–752.
- Dekel, O.; Fischer, F.; and Procaccia, A. D. 2010. Incentive compatible regression learning. *Journal of Computer and System Sciences* 76(8):759–777.
- Dritsoula, L.; Loiseau, P.; and Musacchio, J. 2017. A game-theoretic analysis of adversarial classification. *IEEE Transactions on Information Forensics and Security* 12(12):3094–3109.
- Dwork, C., and Roth, A. 2014. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science* 9(3):211–407.
- Dwork, C. 2006. Differential privacy. In *ICALP*, 1–12.
- Freund, Y., and Schapire, R. E. 1999. Adaptive game playing using multiplicative weights. *Games and Economic Behavior* 29(1-2):79–103.
- Herzog, S.; Shenker, S.; and Estrin, D. 1995. Sharing the cost of multicast trees: An axiomatic analysis. *SIGCOMM Comput. Commun. Rev.* 25(4):315–327.
- Hu, J., and Wellman, M. P. 2003. Nash q-learning for general-sum stochastic games. *Journal of Machine Learning Research* 4:1039–1069.
- Kumar, A., and III, H. D. 2012. Learning task grouping and overlap in multi-task learning. In *ICML*, 1103–1110.
- Liu, W., and Chawla, S. 2009. A game theoretical model for adversarial learning. In *ICDM Workshops*, 25–30.
- Liu, T.-Y.; Chen, W.; and Qin, T. 2015. Mechanism learning with mechanism induced data. In *AAAI*, 4037–4041.
- Mansour, Y.; Mohri, M.; and Rostamizadeh, A. 2009a. Domain adaptation: Learning bounds and algorithms. In *COLT*.
- Mansour, Y.; Mohri, M.; and Rostamizadeh, A. 2009b. Multiple source adaptation and the rényi divergence. In *UAI*, 367–374.
- Meir, R.; Procaccia, A. D.; and Rosenschein, J. S. 2012. Algorithms for strategyproof classification. *Artificial Intelligence* 186:123–156.
- Nash, J. F. 1950. Equilibrium points in n-person games. *National Academy of Sciences of the USA* 36(48-49).
- Nash, J. F. 1951. Non-cooperative games. *Annals of Mathematics* 54(2):286–295.
- Panait, L., and Luke, S. 2005. Cooperative multi-agent learning: The state of the art. *Autonomous Agents and Multi-Agent Systems* 11(3):387–434.
- Peshkin, L.; Kim, K.-E.; Meuleau, N.; and Kaelbling, L. P. 2000. Learning to cooperate via policy search. In *UAI*, 489–496.
- Schuermans, D., and Zinkevich, M. A. 2016. Deep learning games. In *NIPS*, 1678–1686.
- Shalev-Shwartz, S., and Singer, Y. 2007a. Convex repeated games and fenchel duality. In *NIPS*, 1265–1272.
- Shalev-Shwartz, S., and Singer, Y. 2007b. A primal-dual perspective of online learning algorithms. *Machine Learning* 69(2-3):115–142.
- Shapley, L. S. 1953. A value for n-person games. In *Contributions to the Theory of Games*, volume II, 307–317. Princeton University Press.
- Shiva Prasad Kasiviswanathan, Homin K. Lee, K. N. S. R., and Smith, A. 2011. What can we learn privately? *SIAM Journal on Computing* 40(3):793–826.
- Valiant, L. G. 1984. A theory of the learnable. *Communications of the ACM* 27:1134–1142.
- Wang, J.; Kolar, M.; and Srerbo, N. 2016. Distributed multi-task learning. In *AISTATS*, 751–760.