Geometry Correspondence between Empirical and Population Games

Shuang Li, Qiuwei Li, Gongguo Tang, and Michael B. Wakin
Department of Electrical Engineering
Colorado School of Mines
Golden, CO 80401
{shuangli,qiuli,gtang,mwakin}@mines.edu

Abstract

Recent years have seen a blooming application of machine learning models that are trained via artificially constructed differentiable games. Nonlinear optimization, which has a long history and a rich geometric theory, can be viewed as a special differentiable game where there is a single player or all players share a common cost function and control all variables. In this work, we aim at generalizing the ideas and techniques developed for nonlinear optimization to differentiable games. In particular, we develop a generic framework to connect the landscape of an empirical game and its population counterpart. This enables us to characterize the geometric landscape of the empirical risk from the landscape of its population risk, which is typically easier to analyze.

1 Introduction

Game theory provides a mathematical framework to model, reason, and analyze the interactions between several decision-makers, referred to as players, with oftentimes conflicting objectives. A *game* is an interactive situation in which each player's cost or utility depends not only on his/her own actions but also on those taken by the other players. For example, the time it takes for a driver to get home depends not only on the route he/she chooses but also on the other drivers' choices. *Differentiable games* – games with continuous decision variables and differentiable cost functions – have been gradually adopted to model many signal processing, communication, and networking problems in the last two decades (see [1, 16, 7] and references therein). More importantly, the last few years have seen a blooming application of machine learning models that are trained via artificially constructed differentiable games, including generative adversarial networks [6], adversarial training [10], reinforcement learning [17], intrinsic curiosity [13], and imaginative agents [14].

The mathematical formulation of differentiable games is intimately connected to several branches of optimization. The presence of multiple cost functions, one for each player, is tied with multi-objective or vector optimization [4, 8], though they differ in the control of the variables. Partial control of variables relates games to distributed optimization [3], although a common objective function is usually assumed in the latter context. When there is a single player, or all players share a single cost function and control all variables, a differentiable game reduces to nonlinear programming or optimization [12, 2], where a single smooth objective function is optimized in a centralized manner.

Empirical risk minimization is a standard framework in machine learning, where the data generating distribution can only be accessed through data samples. The correspondence of the geometric landscape of empirical risk and that of its population counterpart was first studied in [11] for non-convex losses that is strongly *Morse*, namely, the Hessian of the population risks cannot have zero

Abstract for: Smooth Games Optimization and Machine Learning Workshop (NeurIPS 2019), Vancouver, Canada.

eigenvalues at or near the critical points. This theory was later extended to a more general framework where the strongly Morse assumption is removed [9].

For many differentiable games used in machine learning, including generative adversarial nets and adversarial training, the data generating distribution can also only be accessed through data samples. Motivated by this observation, the work aims at generalizing the ideas and techniques developed for characterizing the geometry of non-convex optimization problems to differentiable games. In particular, we develop a generic framework to connect the landscape of an empirical game and its population counterpart. This enables us to characterize the geometric landscape of the empirical risk from the landscape of its population risk, which is typically easier to analyze.

2 Problem Formulation

Mathematically, a differentiable non-cooperative (also known as strategic or normal-form) game consists of a finite number of K *rational* players, each having a strategy set (or decision or action space) $X_k \subset \mathbb{R}^{p_k}$ and a cost function $g_k(\cdot)$, $k \in [K] \triangleq \{1, \cdots, K\}$. The cost function $g_k(\cdot)$ is a *differentiable* function of the form $g_k : X_1 \times X_2 \times \cdots \times X_K \to \mathbb{R} :\ (\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_K) \to g_k(\boldsymbol{x})$, where $\boldsymbol{x}_k \in X_k$ is the strategy of player k, $\boldsymbol{x} = (\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_K) \in X \subset \mathbb{R}^p$ is the strategy profile, $X = X_1 \times X_2 \times \cdots \times X_K$ is the direct product of the players' strategy sets, and $p = \sum_{k=1}^{K} p_k$. Denote $\boldsymbol{x}_{-k} = (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_{k-1}, \boldsymbol{x}_{k+1}, \ldots, \boldsymbol{x}_k)$ as the strategies taken by all players other than the player k. With some abuse of notation, we also write $\boldsymbol{x} = (\boldsymbol{x}_k, \boldsymbol{x}_{-k})$ and $g_k(\boldsymbol{x}) = g_k(\boldsymbol{x}_k, \boldsymbol{x}_{-k})$.

The *rationality* of the players means that each player's goal is to minimize its own cost function $g_k(\boldsymbol{x}_k, \boldsymbol{x}_{-k})$ by controlling its own strategy $\boldsymbol{x}_k$ while explicitly knowing that $g_k(\boldsymbol{x}_k, \boldsymbol{x}_{-k})$ is also impacted by the other players' strategies $\boldsymbol{x}_{-k}$. The non-cooperative games defined here are *simultaneous* games where all players move simultaneously, or if they do not, the later players are unaware of the earlier players' actions, making them effectively simultaneous. A two-player, zero-sum game has $K = 2$ players where $g_1(\boldsymbol{x}) = -g_2(\boldsymbol{x}) \triangleq g(\boldsymbol{x})$. It becomes a (sequential) minimax game $\min_{\boldsymbol{x}_1} \max_{\boldsymbol{x}_2} g(\boldsymbol{x}_1, \boldsymbol{x}_2)$ when the min-player moves first, and the max-player who moves second can observe the action taken by the first player and adjust its action accordingly.

The data generating distribution in many games used in machine learning can only be accessed through data samples, that is, the cost functions are empirical risks (denoted as $f(\cdot)$) that constructed from finite number of samples. Note that the population counterpart (denoted as $g(\cdot)$) depends on infinite data can be viewed as the expectation of its empirical risk with respect to the randomly generated data samples. Previous works on the geometry of non-convex optimization have shown that the geometry of the population risk is typically easier to analyze [11, 9]. This inspires us to build a general theory that connects the landscape of the empirical game to its population counterpart, which enables us to analyze the geometry of the empirical game from the landscape of its population game. One can refer to Section 4 for a simple example of empirical game and its population counterpart.

3 Main Results

In this section, we present our main results on the connection between the landscape of the empirical game and its population counterpart. Recall that we denote $\boldsymbol{x} = (\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_K) \in X \subset \mathbb{R}^p$ as the strategy profile and $X = X_1 \times X_2 \times \cdots \times X_K$ as the direct product of the players' strategy sets. Here, $\boldsymbol{x}$ can be viewed as the parameter vector that needs to be optimized. Define $\mathcal{B}(l)$ as a compact and connected subset of X with l being a problem-specific parameter. For $k \in [K]$, we denote $f_k(\boldsymbol{x})$ and $g_k(\boldsymbol{x})$ as the empirical risk and the corresponding population risk for player k, respectively.

Differentiability of the cost functions allows the use of calculus tools to characterize local equilibria, a unique feature that is not available to other types of games. Similar to the optimality conditions in optimization, these characterizations form the theoretical and algorithmic foundation for differentiable game optimization. Given a game with cost functions $\{g_k\}_{k\in[K]}$, the *game gradient* is defined as the gradient of the costs with respect to a player's own variables, namely,

$$\mathbf{grad}_g(\boldsymbol{x}) \triangleq \left[(\nabla_1 g_1(\boldsymbol{x}))^\top \ (\nabla_2 g_2(\boldsymbol{x}))^\top \ \cdots \ (\nabla_K g_K(\boldsymbol{x}))^\top \right]^\top \in \mathbb{R}^p,$$

where $\nabla_k g_k(\boldsymbol{x})$ denotes the gradient of $g_k(\boldsymbol{x}) = g_k(\boldsymbol{x}_k, \boldsymbol{x}_{-k})$ with respect to $\boldsymbol{x}_k$. The *game Jacobian* is then the following $(p \times p)$-matrix of second-derivatives

$$\mathbf{J}_g(\boldsymbol{x}) \triangleq \begin{bmatrix} \nabla_{11}^2 g_1 & \nabla_{12}^2 g_1 & \cdots & \nabla_{1K}^2 g_1 \\ \nabla_{21}^2 g_2 & \nabla_{22}^2 g_2 & \cdots & \nabla_{2K}^2 g_2 \\ \vdots & \vdots & \vdots & \vdots \\ \nabla_{K1}^2 g_K & \nabla_{K2}^2 g_K & \cdots & \nabla_{KK}^2 g_K \end{bmatrix} \in \mathbb{R}^{p\times p},$$

where $[\mathbf{J}_g]_{kj}(\boldsymbol{x}) \triangleq \nabla_{kj}^2 g_k(\boldsymbol{x}) \in \mathbb{R}^{p_k \times p_j}$ is a submatrix of the Hessian of $g_k(\boldsymbol{x})$. Note that the game Jacobian is a square and generally *non-symmetric* matrix. A critical point $\boldsymbol{x}$ is a *locally stable equilibrium* if all the eigenvalues of $\mathbf{J}_g(\boldsymbol{x})$ have strictly positive real parts, and is a *locally unstable equilibrium* if at least one of the eigenvalues of $\mathbf{J}_g(\boldsymbol{x})$ have strictly negative real parts.

Optimality conditions characterizing local equilibria are foundational to differential game optimization. For an unconstrained game, i.e., $X = \mathbb{R}^p$, if $\boldsymbol{x}^\star$ is a *local Nash equilibrium*, then

$$\mathbf{grad}_g(\boldsymbol{x}^\star) = \mathbf{0} \text{ and } [\mathbf{J}_g]_{kk}(\boldsymbol{x}^\star) \succeq \mathbf{0},\ k \in [K].$$

For constrained games, the optimality conditions are direct generalizations of the Karush-Kuhn-Tucker (KKT) conditions in nonlinear optimization [15].

Next, we introduce the assumptions needed to build our theory.

Assumption 3.1. *The population risks $\{g_k(\boldsymbol{x})\}_{k=1}^K$ satisfy*

$$\min_{k\in[K]} \min_{i\in[p_k]} |\lambda_i([\mathbf{J}_g]_{kk}(\boldsymbol{x}))| > \eta \quad \textit{and} \quad \min_{j\in[p]} |\mathit{Re}\{\lambda_j([\mathbf{J}_g](\boldsymbol{x}))\}| > \eta$$

in the set $\overline{\mathcal{D}} \triangleq \{\boldsymbol{x} \in \mathcal{B}(l) : \|\boldsymbol{grad}_g(\boldsymbol{x})\|_2 < \epsilon\}$ for certain positive constants ϵ, η. Here, $\lambda_{\min}(\cdot)$ denotes the minimal eigenvalue and $Re\{\cdot\}$ picks the real part of a complex number.

Assumption 3.2. *(Gradient proximity). The game gradients of the empirical risk and population risk satisfy*

$$\sup_{\boldsymbol{x}\in\mathcal{B}(l)} \|\boldsymbol{grad}_f(\boldsymbol{x}) - \boldsymbol{grad}_g(\boldsymbol{x})\|_2 \le \frac{\epsilon}{2}.$$

Assumption 3.3. *(Jacobian proximity). The game Jacobians of the empirical risk and population risk satisfy*

$$\sup_{\boldsymbol{x}\in\mathcal{B}(l)} \|\mathbf{J}_f(\boldsymbol{x}) - \mathbf{J}_g(\boldsymbol{x})\|_2 \le \frac{\eta}{2}.$$

Note that Assumption 3.1 is based on the strongly Morse condition of the population risks$\{g_k(\boldsymbol{x})\}_{k=1}^K$. Assumptions 3.2 and 3.3 can be satisfied with high probability when one has enough data samples. We are now in the position to present our main theory.

Theorem 3.1. *Denote $\{f_k\}_{k\in[K]}$ and $\{g_k\}_{k\in[K]}$ as the empirical risks and the corresponding population risks for the K players in a differentiable game, respectively. Let $\mathcal{D}$ be a maximally* connected *and compact component of the set $\overline{\mathcal{D}}$ with a $\mathcal{C}^2$ boundary $\partial\mathcal{D}$. Under Assumptions 3.1, 3.2, and 3.3 stated above, the following statements hold:*

(a) *If $\{g_k\}_{k\in[K]}$ has one (no, resp.) local Nash equilibrium in $\mathcal{D}$, then $\{f_k\}_{k\in[K]}$ has one (no, resp.) local Nash equilibrium in $\mathcal{D}$.*

(b) *If $\{g_k\}_{k\in[K]}$ has one (no, resp.) non-Nash equilibrium in $\mathcal{D}$, then $\{f_k\}_{k\in[K]}$ also has one (no, resp.) non-Nash equilibrium in $\mathcal{D}$.*

(c) *If $\{g_k\}_{k\in[K]}$ has one (no, resp.) locally stable equilibrium in $\mathcal{D}$, then $\{f_k\}_{k\in[K]}$ has one (no, resp.) locally stable equilibrium in $\mathcal{D}$.*

(d) *If $\{g_k\}_{k\in[K]}$ has one (no, resp.) locally unstable equilibrium in $\mathcal{D}$, then $\{f_k\}_{k\in[K]}$ also has one (no, resp.) locally unstable equilibrium in $\mathcal{D}$.*

The key technical ingredient in the proof of Theorem 3.1 is a classical result in differential topology, which states that the total index of a vector field in a region is fully determined by its behavior near the boundary of that region [15, 5, 9]. Since this fundamental result does not require the vector field to be the gradient of a function, it applies without modification to the vector field of the game gradient. The above theorem indicates that the correspondence between locally stable and unstable equilibria should be relatively straightforward, as they are similar to local minima and strict saddles in optimization.

4 Simulations

In this section, we conduct numerical experiments to further support our theory. For simplicity, we test our theory on a two-player, zero-sum game ($K = 2$) with empirical risks being $f_1(x_1, x_2) = -f_2(x_1, x_2) \triangleq f(x_1, x_2)$. Denote $\boldsymbol{q}^\star = [1/4; -1; 1/2; -1/4] \in \mathbb{R}^4$ as the true coefficients of the model. Assume that we only have access to its noisy samples $\boldsymbol{q}_m = \boldsymbol{q}^\star + \boldsymbol{e}_m,\ m \in [M]$, where $\boldsymbol{e}_m$ denotes the observation error with entries following the Gaussian distribution $\mathcal{N}(0, 1)$. We then construct the cost function as

$$f(x_1, x_2) \triangleq [x_1^4\ \ x_1^2\ \ x_1^2x_2^2\ \ x_2^2]\frac{1}{M}\sum_{m=1}^{M}\boldsymbol{q}_m = [x_1^4\ \ x_1^2\ \ x_1^2x_2^2\ \ x_2^2](\boldsymbol{q}^\star + \frac{1}{M}\sum_{m=1}^{M}\boldsymbol{e}_m).$$

Note that the corresponding population risk is then given as

$$g(x_1, x_2) = \mathbb{E}f(x_1, x_2) = [x_1^4\ \ x_1^2\ \ x_1^2x_2^2\ \ x_2^2]\boldsymbol{q}^\star = \frac{1}{4}x_1^4 - x_1^2 + \frac{1}{2}x_1^2x_2^2 - \frac{1}{4}x_2^4,$$

which has seven critical points: $(x_1, x_2) = (0, 0),\ (\pm\sqrt{2}, 0),\ (\pm 1, \pm 1)$. Among these critical points, $(0, 0),\ (\pm\sqrt{2}, 0)$ are locally unstable equilibria and $(\pm 1, \pm 1)$ are locally stable equilibria. We present the landscapes of the population risk and two realizations of the empirical risk with $M = 100$ and $M = 200$ in Figure 1, which indicates a direct correspondence between the locally stable equilibria and locally unstable equilibria of the empirical game and its population counterpart. In addition, for this particular two-player, zero-sum game, one can verify that the four locally stable equilibria $(x_1, x_2) = (\pm 1, \pm 1)$ are also local Nash equilibria.

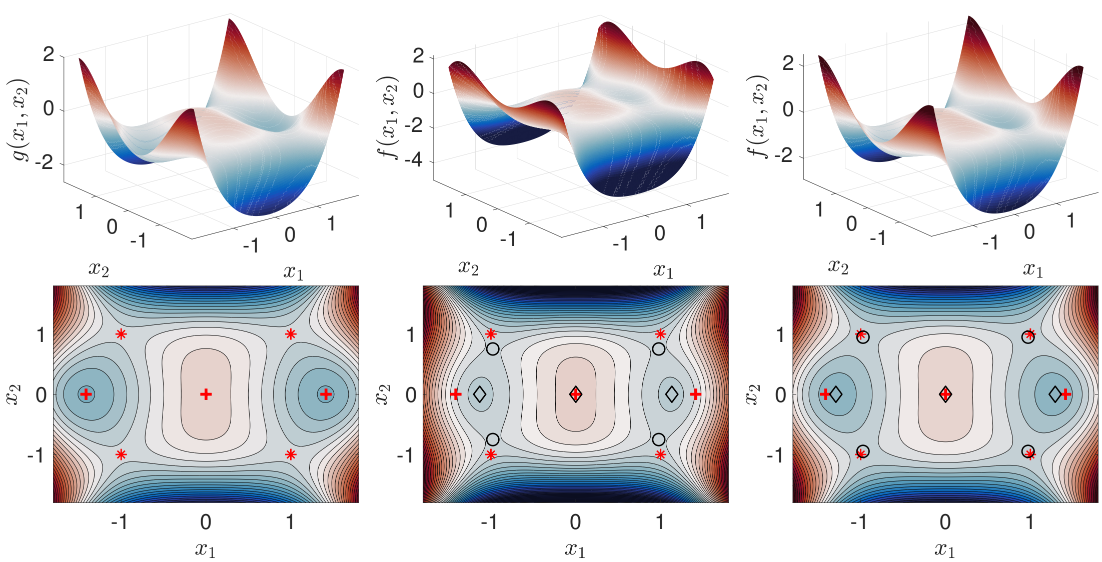

(a) (b) $M = 100$ (c) $M = 200$

Figure 1: The landscape of a two-player, zero-sum game: (a) Population risk. (b, c) Two realizations of the empirical risk. We use the red star and cross to denote the locally stable equilibria and locally unstable equilibria of the population risk, respectively, and use the black circle and diamond to denote the locally stable equilibria and locally unstable equilibria of the empirical risk, respectively.

5 Conclusions

In this work, we built a general theory that connects the geometric landscape of the empirical game to its population counterpart by generalizing the ideas and techniques developed for characterizing the geometry of non-convex optimization problems to differentiable games. In particular, we developed a theory that establishes a correspondence of the empirical game and its population counterpart's critical points. This one-to-one correspondence between the landscape of empirical risk and that of the population risk makes it possible for us to characterize the geometric landscape of the empirical risk from the landscape of its population counterpart.

References

[1] G. Bacci, S. Lasaulce, W. Saad, and L. Sanguinetti. Game theory for networks: A tutorial on game-theoretic tools for emerging signal processing applications. *IEEE Signal Processing Magazine*, 33(1):94–119, 2015.

[2] D. P. Bertsekas. Nonlinear programming. *Journal of the Operational Research Society*, 48(3):334–334, 1997.

[3] S. Boyd, N. Parikh, E. Chu, B. Peleato, J. Eckstein, et al. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning*, 3(1):1–122, 2011.

[4] K. Deb. *Multi-objective optimization using evolutionary algorithms*, volume 16. John Wiley & Sons, 2001.

[5] B. A. Dubrovin, A. T. Fomenko, and S. P. Novikov. *Modern geometry—methods and applications: Part II: The geometry and topology of manifolds*, volume 104. Springer Science & Business Media, 2012.

[6] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2672–2680, 2014.

[7] Z. Han, D. Niyato, W. Saad, T. Basar, and A. Hjorungnes. *Game theory in wireless and communication networks: theory, models, and applications*. Cambridge University Press, 2012.

[8] J. Jahn. *Vector optimization*. Springer, 2009.

[9] S. Li, G. Tang, and M. B. Wakin. The landscape of non-convex empirical risk with degenerate population risk. In *Advances in Neural Information Processing Systems*, 2019.

[10] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.

[11] S. Mei, Y. Bai, and A. Montanari. The landscape of empirical risk for non-convex losses. *arXiv preprint arXiv:1607.06534*, 2016.

[12] J. Nocedal and S. Wright. *Numerical optimization*. Springer Science & Business Media, 2006.

[13] D. Pathak, P. Agrawal, A. A. Efros, and T. Darrell. Curiosity-driven exploration by self-supervised prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 16–17, 2017.

[14] S. Racanière, T. Weber, D. Reichert, L. Buesing, A. Guez, D. J. Rezende, A. P. Badia, O. Vinyals, N. Heess, Y. Li, et al. Imagination-augmented agents for deep reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 5690–5701, 2017.

[15] J. B. Rosen. Existence and uniqueness of equilibrium points for concave n-person games. 1964.

[16] G. Scutari, D. P. Palomar, F. Facchinei, and J.-s. Pang. Convex optimization, game theory, and variational inequality theory. *IEEE Signal Processing Magazine*, 27(3):35–49, 2010.

[17] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484, 2016.