

Game theory and artificial intelligence

Advisor (s): Sébastien Gerchinovitz (IRT Saint-Exupéry, UT3) and Jérôme Renault (Université Toulouse 1 Capitole - TSE) - jerome.renault@tse-fr.eu
sebastien.gerchinovitz@irt-saintexupery.com

<https://sites.google.com/site/jrenaultsite>

Net salary: 2096€ per month with some teaching (64 hours per year on average)

Duration: 36 months

PRACTICAL DETAILS

This 3-year PhD project is funded by the new Artificial and Natural Intelligence Toulouse Institute (ANITI). It will be part of the chair Game Theory and Artificial Intelligence. We are looking for an excellent candidate with a strong background in mathematics (e.g. in game theory, optimization, analysis, probability and statistics) or theoretical computer science. The PhD can start at any time from September 2019 to September 2020, with a monthly net salary of 2097 euros and 64 hours of teaching every year.

RESEARCH TOPIC

Generative Adversarial Networks (GANs) are a class of unsupervised machine learning techniques to estimate a distribution from high-dimensional data and to sample elements that mimic the observations (Goodfellow et al., 2014). They use a zero-sum dynamic game between two neural networks: a generator, which generates new “fake” data instances, and a discriminator, which evaluates true and fake instances. While GANs are receiving a huge interest, it is important to improve their training and to provide theoretical guarantees on their convergence and robustness.

Though GANs are explicitly defined as a game between two networks, most algorithms that have been studied so far (mostly empirically) do not use this fact. It is actually frequent to see two stochastic gradient descent algorithms used independently for the generator and the discriminator, with no proved guarantees on why this should lead to a good joint solution of the min-max problem.

There have been recent attempts to address the GAN training task as a min-max problem per se, through game-theoretic/online-learning approaches; see, e.g., Ge et al. (2018); Arjovsky et al. (2017); Grnarova et al. (2018); Hsieh et al. (2019). We would like to further explore these directions, with a careful attention on computational issues. We would like to investigate how some algorithms that have been originally designed for games could prove useful for GAN training. These algorithms may come from non-cooperative game theory (beyond fictitious play) or from adversarial online learning (no-regret algorithms), and will need to be largely adapted to the GAN context. Convergence bounds may be combined with approximation and generalization guarantees for neural networks.

In addition, payoffs in GANs may be modified, e.g., they could be defined through new probability metrics designed for partially observable Markov decision processes and dynamic games with incomplete information. More generally, it is worthwhile to try to go beyond the 2-player zero-sum game setup and design new games to learn and mimic distributions with high-dimensional data, with improved empirical and theoretical properties.

References

- Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In Proceedings of the 34th International Conference on Machine Learning, pages 214–223, 2017.
- G  rard Biau, Beno  t Cadre, Maxime Sangnier, and Ugo Tanielian. Some theoretical properties of gans. arXiv preprint arXiv:1803.07819, 2018.
- Hao Ge, Yin Xia, Xu Chen, Randall Berry, and Ying Wu. Fictitious gan: Training gans with historical models. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, Computer Vision – ECCV 2018, pages 122–137, Cham, 2018. Springer International Publishing.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, Advances in Neural Information Processing Systems 27, pages 2672–2680. Curran Associates, Inc., 2014.
- Paulina Grnarova, Kfir Y. Levy, Aurelien Lucchi, Thomas Hofmann, and Andreas Krause. An online learning approach to generative adversarial networks. In Proceedings of International Conference on Learning Representations (ICLR’18), 2018.
- Ya-Ping Hsieh, Chen Liu, and Volkan Cevher. Finding mixed nash equilibria of generative adversarial networks. In Proceedings of International Conference on Learning Representations (ICLR’19), 2019..

APPLICATION PROCEDURE

Formal applications should include detailed cv including recent academic marks and ranks, up to three recommendation letters stating your ability for research, a short research statement. Samples of published research by the candidate if any will be a plus.
> applications should be sent by email to: advisor_email
More information about ANITI: <https://aniti.univ-toulouse.fr/>

