ANITI
ARTIFICIAL & NATURAL INTELLIGENCE
TOULOUSE INSTITUTE

Chair: Game Theory and Artificial Intelligence

J. Renault
F. Gensbittel
S. Gerchinovitz

September 2019

Université Fédérale
Toulouse Midi-Pyrénées

Chair members

Jérôme Renault

(Chair Holder)
PR maths UT1-TSE

Game Theory, Stochastic games, Economic theory, POMDP, Optimization, Discrete processes, Discrete maths

Fabien Gensbittel

(Associate Researcher)
MCF maths UT1 -TSE

Stochastic games, Dynkin games, Differential games, Incomplete information, Bayesian persuasion

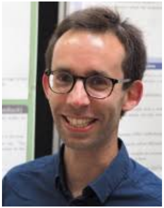

Sébastien Gerchinovitz

(Associate Researcher)
IRT-DEEL Researcher

Online learning, Adversarial bandits, Nonparametric statistics, Optimization

Post-doc 1 (starting Nov 1, 2019)

Tommaso Cesari

PhD in Computer Science at U. of Milan (with N. Cesa-Bianchi)

Optimization, online learning on graphs, regret bounds, dynamic pricing, repeated auctions.

Context

Theoretical chair.

Game theory studies strategic interactions between several agents : auctions, chess, routing and congestion, oligopolies, elections...

Goal: to develop links between
Game Theory and Artificial Intelligence.

Game Theory and AI: summary

Part 1. GANs : Generative Adversarial Networks, where game theory is used to improve "traditional" AI tasks.
Invented in 2014 by I. Goodfellow, "the coolest idea in the last 20 years in machine learning", ac. to Yann Le Cun in 2016.

Part 2. Complex strategic interactions (between humans, between AI systems or between a mix of humans and AI systems).
Think of repeated auctions for Google ads, autonomous vehicles, complex markets.
Develop theoretical properties (appropriate notions of solutions, existence, computation) in these games.

Remark: we will *not* study AI software for playing games (Alphago, Alphazero).

1. Generative Adversarial Networks

Goal: estimate a distribution from high-dimensional data and sample elements that mimic the observations (cats, cars, old portraits...)

Idea: Introduce a game between 2 neural networks

- the Generator trying to generate real-looking images,
- the Discriminator trying to distinguish between real and fake images.

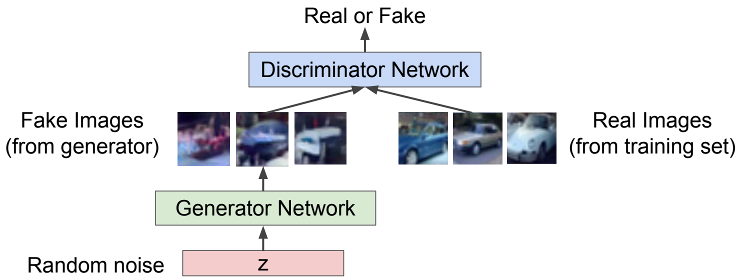

Fake and real images copyright Emily Denton et al. 2015.

Minmax objective function:

$$\min_{\theta_g} \max_{\theta_d} \left(\mathop{\mathbb{E}}_{x\sim p_{data}} \log D_{\theta_d}(x) + \mathop{\mathbb{E}}_{z\sim p(z)} \log(1 - D_{\theta_d}(G_{\theta_g}(z))) \right)$$

with $D_{\theta_d}(x) \in [0,1]$ probability assigned by D that x is true.

Discriminator chooses θ_d and wants $D_{\theta_d}(x)$ close to 1 if x is real, and close to 0 if x is fake.
Generator chooses θ_g and wants $D_{\theta_d}(x)$ close to 1 it x is fake.

Optimal strategy (if available) for the generator: chooses θ_g such that $G_{\theta_g}(z) \sim p_{data}$.

Training GANs: typically alternate between stochastic gradient ascent on discriminator and descent on generator.
(Many variants for payoffs, training)

Works very well in practice. (sometimes too well: deepfake)

 

Several important issues:

Understand why the game formulation is so efficient. Convergence bounds of the algorithms, combined with approximation and generalization guarantees for neural networks. Investigate other algorithms used in games. Imagine/invent other games with interesting algorithmic properties

2. Complex strategic interactions

Many interesting situations deserve investigation. In particular:

• 2.a) Games with differentiable (and ReLU) payoff functions. In particular when agents choose parameters of their own neural networks, or when they are restricted to use low complexity strategies.

• 2.b) Solutions of general complex environments such as stochastic games, dynamic games with incomplete information and/or with signals or POMDP, real-time games ?

• 2.c) Hybrid games where some players are restricted to a class of algorithms while others are not.
Need for safe and robust design of strategies/algorithms, possibly with a loss of efficiency.

- 2.d) Continuous Bandit Optimization: minimization of an unknown function, estimation of level sets. Application to certification of neural networks.

- 2.e) Cooperation and Competition in bandits with congestion

- 2.f) Complex interactions between humans : Matching Bandits (or Bandit Markets à la M. Jordan)

Possible interactions:

(Stochastic) gradient descent, and neural networks in GANs : Optimization chairs.

Matching bandits: economics chair

+ ?

PhD proposal 1:

Game theory for Generative Adversarial Networks.
(Advisors S.Gerchinovitz, J. Renault)

PhD proposal 2:

Game Theory and Artificial Intelligence.
(Advisors F. Gensbittel, J. Renault)

+ Post-doc 2