Study Group on FM and AI

When: Once a month?

Aim: understand AI methods (learning…), from FM point of view. Share what we know, learn and read.

For: - Use AI methods in our FM works (more efficiency…)
- Develop FM techniques to answer some problems in AI (e.g.: trustworthy learning)
- Understand for what AI methods work best, and for what FM methods work best
(e.g.: AI for classification, MF for analysing dynamical behaviors)

Today: Replay of Summit on Machine Learning Meets Formal Methods, Oxford 2018.
https://easychair.org/smart-slide/conference/FBG9 (see also talk by Platzer https://easychair.org/smart-slide/slide/jTdT#)
Talk of Marta Kwiatkowska Safety verification for deep neural networks with provable guarantees
Based on:
CAV 2017 https://arxiv.org/pdf/1610.06940.pdf
IJCAI 2018 https://arxiv.org/pdf/1805.02242.pdf
TACAS 2018 https://arxiv.org/pdf/1710.07859.pdf

Demystifying AI

Talk by Kwiatkowska shows us DNN is not entirely satisfactory even for what it does best (classification).

Use in videogame is quite impressive, but:

https://blog.openai.com/openai-five/

Solo skills (Reinforcement learning):
Beat every human
But…
Learn from 30.000 years of plays

	OPENAI 1V1 BOT	OPENAI FIVE
CPUs	60,000 CPU cores on Azure	128,000 preemptible CPU cores on GCP
GPUs	256 K80 GPUs on Azure	256 P100 GPUs on GCP
Experience collected	~300 years per day	~180 years per day (~900 years per day counting each hero separately)
Size of observation	~3.3 kB	~36.8 kB
Observations per second of gameplay	10	7.5
Batch size	8,388,608 observations	1,048,576 observations
Batches per minute	~20	~60

Demystifying AI

Talk by Kwiatkowska shows us DNN is not entirely satisfactory even for what it does best (classification).

Use in videogame is quite impressive, but:

https://blog.openai.com/openai-five/

	OPENAI 1V1 BOT	OPENAI FIVE
CPUs	60,000 CPU cores on Azure	128,000 preemptible CPU cores on GCP
GPUs	256 K80 GPUs on Azure	256 P100 GPUs on GCP
Experience collected	~300 years per day	~180 years per day (~900 years per day counting each hero separately)
Size of observation	~3.3 kB	~36.8 kB
Observations per second of gameplay	10	7.5
Batch size	8,388,608 observations	1,048,576 observations
Batches per minute	~20	~60

Teamplay (Reinforcement learning):
Learn from 100.000 years of plays
Good results, but still not at the top

https://www.youtube.com/watch?v=nGhkCuQloXI

Showmatch 1 [edit]

Place	Team
W	paiN Gaming
L	OpenAI Five

Showmatch 2 [edit]

Place	Team
W	Big God
L	OpenAI Five

Next 68NQRT by Hugo Bazille on 11 October : Learning trustworthy Markov Chains.

Topics to be presented

- Planning (logica)
- Reinforcement learning
- Process discovery
- SAT/SMT solvers (heuristics)
- Learning automata (Angluin / passive learning…), Testing – learning timed automata.
- Distance between automata =>
- Supervisory control and learning.
- Understanding backpropagation in DNN.

November: different approaches for learning models

December: Eric?

January – logica on planing