>>> Quantum machine learning for hackers: a primier on Projective Simulation

>>> EndSummerCamp 2016 - Venice

Name: Alessandro 'Scinawa' Luongo†

Date: July 10, 2018

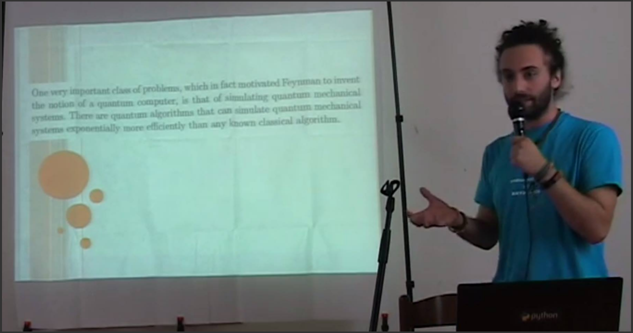

Figure: Scinawa @ EndSummerCamp 2014

†alessandro@luongo.pro - Twitter: @scinawa

>>> Table of Content

>>> Machine Learning ... wut?

* Machine learning algorithms learn a desired input-output relation from examples in order to interpret new inputs [Schuld et al., 2015]
* Machine learning is like money laundering for bias [a dude on Twitter, 2016]

>>> Machine Learning ... wut?

* Machine learning algorithms learn a desired input-output relation from examples in order to interpret new inputs [Schuld et al., 2015]
* Machine learning is like money laundering for bias [a dude on Twitter, 2016]

1. Unsupervised $\{x_1, ...x_n,\}$
2. Supervised $\{(x_1, y_1)...(x_n, y_n)\}$
3. Reinforcement Learning: (get_action + get_reward)

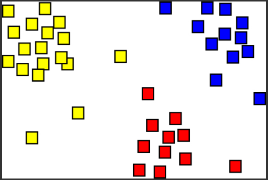

Figure: [Wikipedia, 2016]

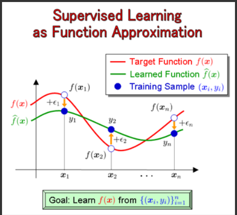

Figure: [tok,]

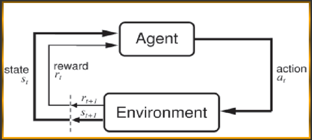

Figure: [Ahmed et al., 2002]

>>> Quantum mechanics: bugs or features?

* Qubit = 2 level system: $|0\rangle = [1, 0], |1\rangle = [0, 1]$
* Physical states are (unitary) vectors $\in \mathbb{H}^n \simeq \mathbb{C}^{2^n}$
* Evolution as (unitary) matrices
* CCNOT like NAND

Figure: [Wikipedia, 2016]

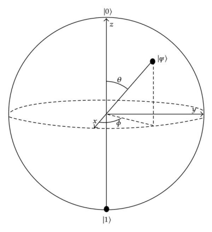

Figure: [Wikipedia, 2016]

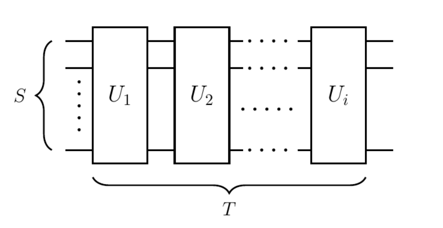

>>> Bugs or features?

* Computation is reversible and linear. Measurements are irreversible.
* Entanglement: spooky action at a distance
* No-cloning: you cannot copy (pure) quantum information.
* Superposition: $|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle$
* Interference: particles behave like waves and interfere as such

>>> Bugs or features?

* Computation is reversible and linear. Measurements are irreversible.
* Entanglement: spooky action at a distance
* No-cloning: you cannot copy (pure) quantum information.
* Superposition: $|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle$
* Interference: particles behave like waves and interfere as such

* Grover's algorithm finds the input x to a function $f(x)$ that produces a particular output value y, using just $O(N^{1/2})$ evaluations of the function, where N is the size of the function's domain. [Wikipedia, 2016]
* HHL algorithm for matrix inverse

>>> Why it is important to ML?

* Speedup (which trade-offs?) [Wittek, 2014]
* Storage capacity is of interest. [Wittek, 2014]
* New models for data [Wiebe, 2016]
* Efficient training with fewer approximations [Wiebe, 2016]
* What is learning?

>>> List of qML algorithm

Table: Main approaches to ML: [Wittek, 2014]

Algorithm	Use Grover	Speedup
K-Medians	Y	Quadratic
Hierarchical clustering	Y	Quadratic
K-means	Opt.	Exponential
PCA	N	Exponential
Neural Networks	Y	-
Support Vector Machines	Y	Quadratic
Support Vector Machines	N	Exponential
Nearest neighbors	Yes	Quadratic
Regression	N	-
Projective Sim.	Y	Quadratic

* Exhibit 1: Quantum CS vs Classical CS

* Exhibit 1: Quantum CS vs Classical CS
* Exhibit 2: Need for specialization in cybersecurity

* Exhibit 1: Quantum CS vs Classical CS
* Exhibit 2: Need for specialization in cybersecurity
* Exhibit 3 ...

>>> Why it is important to us?

* AI^2: Training a big data machine to defend [Veeramachaneni and Arnaldo,]
* This AI Will Craft Tweets That You' ll Never Know Are Spam [Simonite, 2016]
* ``Machine-Learning Algorithm Combs the Darknet for Zero Day Exploits, and Finds Them'' - Darknet and Deepnet Mining for Proactive Cybersecurity Threat Intelligence [Nunes et al., 2016]
* APPLIED MACHINE LEARNING FOR DATA EXFIL AND OTHER FUN TOPICS [Matt Wolff, 2016]
* Cyber Grand Challange (https://www.cybergrandchallenge.com/)

Dan Geer @ RSA 2015 : The future of Cybersecurity

>>> PoC || GTFO

* https://bitbucket.org/mroystein/projectivesimulation
* Quipper: quantum language [Green et al., 2013] (HELP!)
* http://scikit-learn.org/
* Markov exfiltration:
  https://github.com/bwall/markovobfuscate
  [Matt Wolff, 2016]

Figure: [Bjerland, 2015]

>>> The fuck

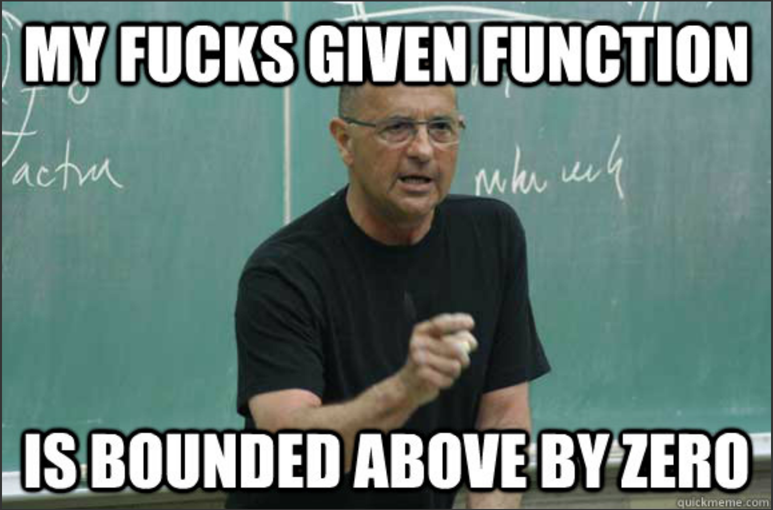
MY FUCKS GIVEN FUNCTION
IS BOUNDED ABOVE BY ZERO
quickmeme.com

>>> Q&A: quantum computers vs. reality

* I'm not a physicist <3

>>> Q&A: quantum computers vs. reality

* I'm not a physicist <3
* 1 Billion euro investment in 2018 (Europe's Quantum Manifesto)

>>> Q&A: quantum computers vs. reality

* I'm not a physicist <3
* 1 Billion euro investment in 2018 (Europe's Quantum Manifesto)
* 11 different ways of implementing qubits.

>>> Q&A: quantum computers vs. reality

* I'm not a physicist <3
* 1 Billion euro investment in 2018 (Europe's Quantum Manifesto)
* 11 different ways of implementing qubits.
* We demonstrate two-qubit and single-qubit logic gates with respective fidelities 99.9(1)% and 99.9934(3)%, significantly above the $\equiv 99\%$ minimum threshold level required for fault-tolerant quantum computation, [...] in a room-temperature trap. (4 Aug.)

>>> Q&A: quantum computers vs. reality

* I'm not a physicist <3
* 1 Billion euro investment in 2018 (Europe's Quantum Manifesto)
* 11 different ways of implementing qubits.
* We demonstrate two-qubit and single-qubit logic gates with respective fidelities 99.9(1)% and 99.9934(3)%, significantly above the $\equiv 99\%$ minimum threshold level required for fault-tolerant quantum computation, [...] in a room-temperature trap. (4 Aug.)
* In addition, quantum computers restricted to domain problems are becoming feasible. [Wittek, 2014]

>>> Q&A: quantum computers vs. reality

* I'm not a physicist <3
* 1 Billion euro investment in 2018 (Europe's Quantum Manifesto)
* 11 different ways of implementing qubits.
* We demonstrate two-qubit and single-qubit logic gates with respective fidelities 99.9(1)% and 99.9934(3)%, significantly above the $\equiv 99\%$ minimum threshold level required for fault-tolerant quantum computation, [...] in a room-temperature trap. (4 Aug.)
* In addition, quantum computers restricted to domain problems are becoming feasible. [Wittek, 2014]
* 15 (2001) ... 143 (2012) ... 56153 (2014) ...

>>> End

E questo è quanto [Plank]'

>>> Projective Simulation: [Briegel and De las Cuevas, 2012]

Figure: Aplysia (Wikipedia)

>>> Episodic and Compositional Memory

* percepts: possible inputs of the algorithm.
  $s = (s_1, s_2, ..., s_N) \in S$, where S is the percept space $S_1 \times ... \times S_n$. Each $s_i = 1, ..., |S_i|$.
* clips: the fundamental unit of the episodic memory. Percept clips are stimulated by percepts, action clips a, if stimulated, trigger an action .
* edges : directed arc between clips. and they contain data useful to the execution of the algorithm.
* emotions: piece of data (called emotions tag) attached to the edge. $e = (e_1, e_2, ..., e_k)$ in the emotional space $E \equiv E_1 \times ...E_k = E, e_k = 1, ...|E_k|$.

>>> Episodic Compositional Memory

Figure: [Briegel and De las Cuevas, 2012]

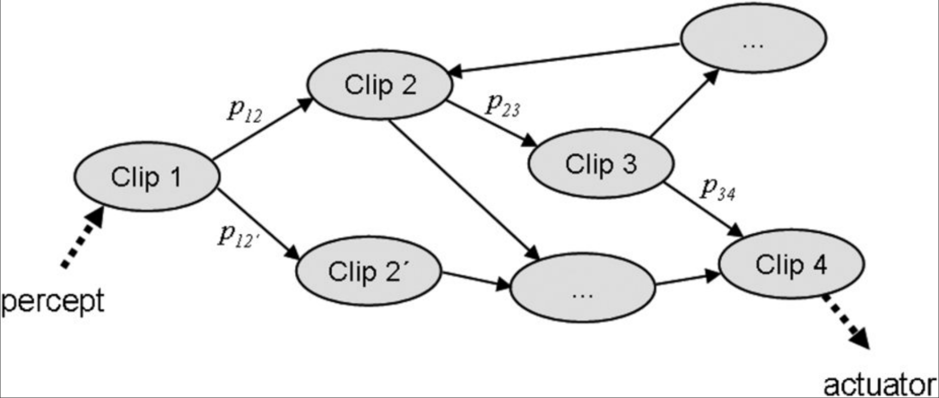

>>> Episodic and Compositional Memory

```
while(1)
    initial_clipt=get_initial_clip(percept)
    (path,action)=get_path_action(initial_clip) #random walk
    for i=0 to R
      if (emotions_ok(path))
        reward=get_reward(action) # exec action
        updates_ecm_memory(reward)
        break
      (path,action)=get_action(initial_clip) #random walk
    # default
    reward=get_reward(action) #exec last action
    updates_ecm_memory(reward)
```

>>> Invasion Game [Briegel and De las Cuevas, 2012]

Figure: [Briegel and De las Cuevas, 2012]

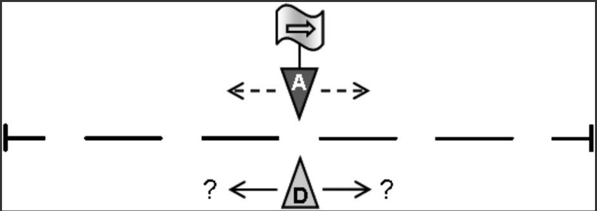

>>> Reflection

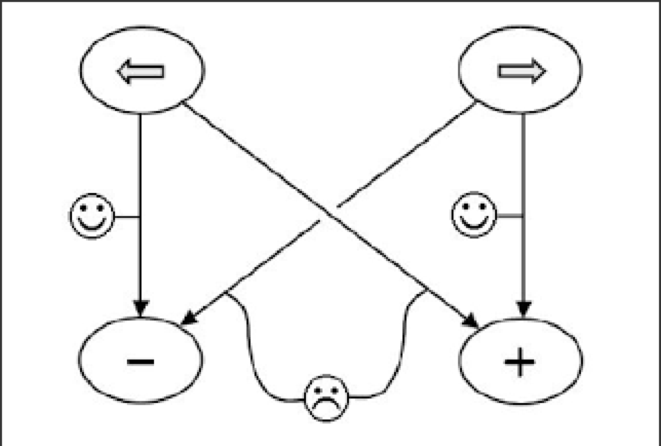

Figure: EC Memory for Invasion game: [Briegel and De las Cuevas, 2012]

>>> Reflection's efficiency

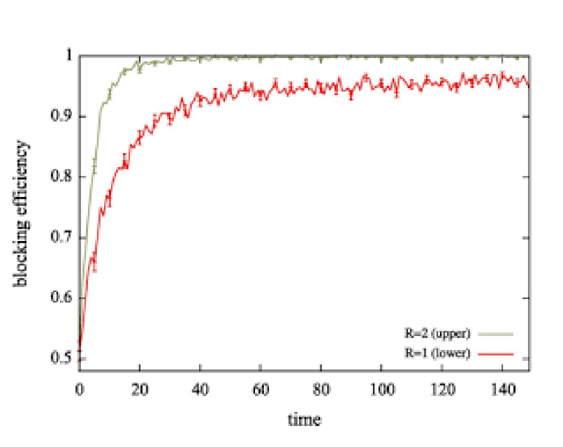

Figure: Reflection r0x: [Briegel and De las Cuevas, 2012]

$$E(t) = \sum_s P^t_{ok}(a|s) P^t(s)$$

>>> Quantum Speedups in PS

* Because of reversibility, you lose information on the arc's direction.
* Grover's algorithm can be used in order to get a quadratic speedup of the ``reflection'' process. [Paparo et al., 2014]
* Exponential speedup if EC Memory is a specific kind of graph []

>>> Why I like it

1. the more you think, the better (often)
2. the more you exercise, the fastest (direct connection percept clip-action clip)
3. Associative learning (creativity)
4. It is biologically inspired and quantum powered
5. Clips composition

Ahmed, M. N., Yamany, S. M., Mohamed, N., Farag, A. A., and Moriarty, T. (2002).
A modified fuzzy c-means algorithm for bias field estimation and segmentation of mri data.
IEEE transactions on medical imaging, 21(3):193--199.

Bjerland, O. F. (2015).
Projective simulation compared to reinformcement learning.

Briegel, H. J. and De las Cuevas, G. (2012).
Projective simulation for artificial intelligence.
Scientific reports, 2.

Green, A. S., Lumsdaine, P. L., Ross, N. J., Selinger, P., and Valiron, B. (2013).
An introduction to quantum programming in quipper.
In *International Conference on Reversible Computation*, pages 110--124. Springer.

Matt Wolff, Brian Wallace, X. Z. (August 4, 2016).
Applied machine learning for data exfil and other fun topics.

Nunes, E., Diab, A., Gunn, A., Marin, E., Mishra, V., Paliath, V., Robertson, J., Shakarian, J., Thart, A., and Shakarian, P. (2016).
Darknet and Deepnet Mining for Proactive Cybersecurity Threat Intelligence.
ArXiv e-prints.

Paparo, G. D., Dunjko, V., Makmal, A., Martin-Delgado, M. A., and Briegel, H. J. (2014).
Quantum speedup for active learning agents.
Physical Review X, 4(3):031002.

Schuld, M., Sinayskiy, I., and Petruccione, F. (2015).
An introduction to quantum machine learning.
Contemporary Physics, 56(2):172--185.

Simonite, T. (August 4, 2016).
This ai will craft tweets that you' ll never know are spam.

Veeramachaneni, K. and Arnaldo, I.
Ai2: Training a big data machine to defend.

Wiebe, N. (2016).
Quantum computing and ai tutorial.

Wikipedia (2016).
Bloch sphere.

Wittek, P. (2014).
Quantum machine learning: what quantum computing means to data mining.
Academic Press.