Overview of Ellis QPhML

Bert Kappen, Riccardo Zecchina

Recent AI revolution

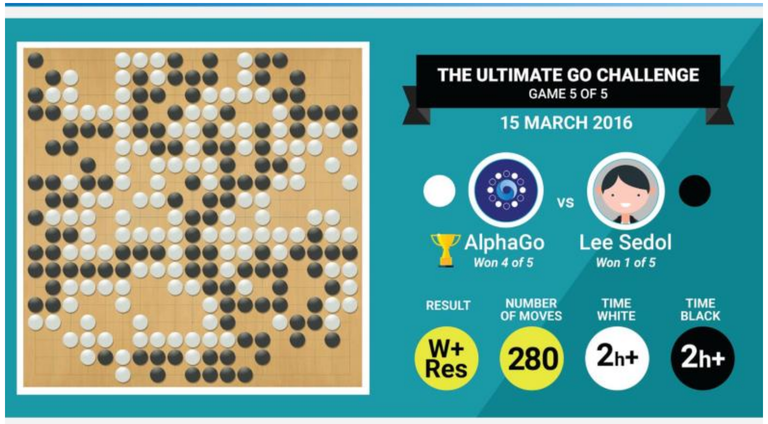

Deep learning of Atari games

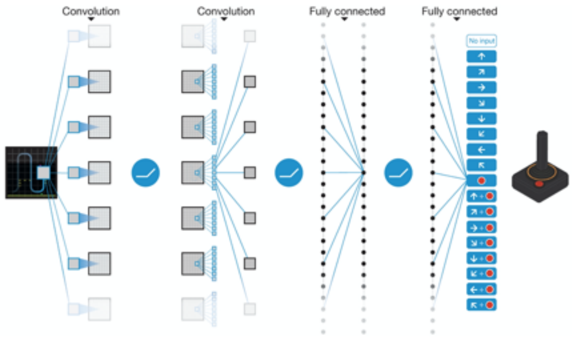

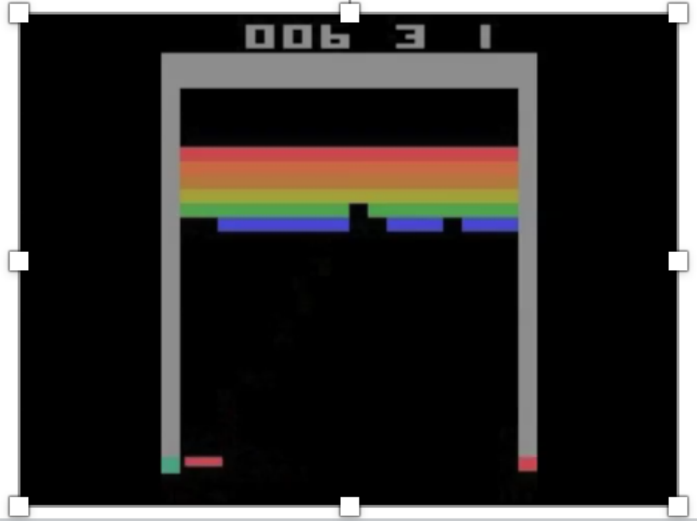

Understanding images

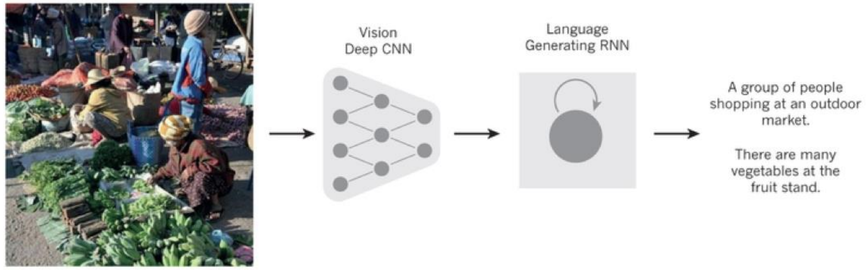

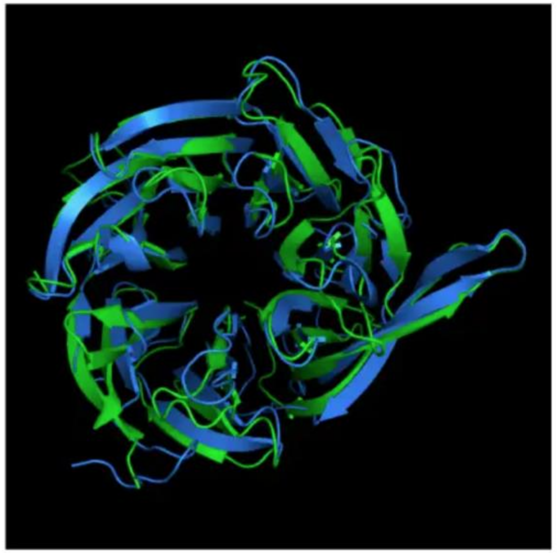

ellis

European Laboratory for Learning and Intelligent Systems

Our Mission

We are at a crossroads where

1. **Machine learning is at the heart of a technological and societal artificial intelligence revolution** involving multiple sister disciplines, with large implications for the future competitiveness of Europe.
2. **Europe is not keeping up:** many of the top labs, as well as many of the top places to do a PhD, are located in North America; moreover, AI investments in China and North America are significantly larger than in Europe.
3. **the distinction between academic research and industrial labs is vanishing**, with a significant part of the basic research now being done in industry (with substantial research freedom, and higher salaries), rapid commercialization of results, and academic institutions worldwide struggling to retain their best scientists (with negative implications not only for research but also for the education of future talent). This further weakens Europe since most of the companies doing top research in this field are controlled from the US (or China) – many European companies whose future business crucially depends on AI are not perceived as competitive.

As an important measure to address these points we propose to found a **European Lab for Learning and Intelligent Systems** (ELLIS), involving the very best European academics while working together closely with basic researchers from industry.

The mission of ELLIS is to benefit Europe in two ways:

1. We want the best basic research to be performed in Europe, to enable Europe to shape how machine learning and modern AI change the world.
2. We want to have economic impact and create jobs in Europe, and believe this is achieved by outstanding and free basic research, independent of industry interests.

While the first batch of Programs was reviewed by CIFAR LMB, subsequent applications are being reviewed by existing ELLIS fellows. For submissions and questions please email **daniela.diaconu@ellis.eu**.

At this point, we have a set of 13 ELLIS Programs:

- ELLIS Health
- ELLIS Robot Learning: Closing the Reality Gap!
- Geometric Deep Learning
- Human-centric Machine Learning
- Interactive Learning and Interventional Representations
- Machine Learning and Computer Vision
- Machine Learning for Earth and Climate Sciences
- Natural Intelligence
- Natural Language Processing
- Quantum and Physics Based Machine Learning
- Robust Machine Learning
- Semantic, Symbolic and Interpretable Machine Learning
- Theory, Algorithms and Computations of Modern Learning Systems

ELLIS PhD & Postdoc Program

The ELLIS PhD & Postdoc Program supports excellent young researchers by connecting them to leading researchers across Europe and offering a variety of networking and training activities, including summer schools and workshops. ELLIS PhDs and postdocs conduct cutting-edge curiosity-driven research in machine learning or a related research area with the goal of publishing in top-tier conferences in the field.

ELLIS PhD Program - Call for applications - Deadline for applications: November 15, 2021

Tracks

There are two tracks within the ELLIS PhD & Postdoc Program: the academic track and the industry track. These tracks have separate requirements for admission and criteria for activity during the appointment, but otherwise offer the same benefits, network and resources to the applicant.

Academic track

PhD students and postdocs in the academic track strive for international collaboration as they partner with two European academic institutions in their research. These candidates are supervised by one **ELLIS fellow/scholar** or unit faculty and one **ELLIS fellow/scholar/member** from different European countries, and they visit the exchange institution for min. 6 months (the partitioning of this time is flexible). Normally, the exchange is partially sponsored by the exchange institution; ELLIS PhD students and postdocs are also eligible to apply for the **ELISE mobility grant**.

Industry track

The industry track is open to PhD and postdoc candidates that will be part of a collaboration between an academic institution and an industry partner, and will spend time conducting research at the industry partner during their PhD or postdoc. The candidate will spend a minimum of 50% of their time at the academic institution, and at least 6 months (cumulative) with the industry partner.

For this track, both advisors may be located in the same country. One advisor will represent the academic institution and the other the industry partner. Both the academic and industry advisor must be **ELLIS members** and at least one of them a **fellow, scholar** or unit faculty.

Past Events

Second NLP ELLIS Workshops

30 Jul 2021 - 30 Jul 2021

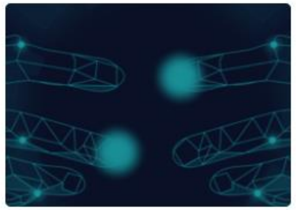

Artificial Intelligence and Robotics in the perspective of social challenges

28 Jul 2021 - 28 Jul 2021

ELLIS workshop on Causethical ML

26 Jul 2021 - 26 Jul 2021

GSI 2021 - LEARNING GEOMETRIC STRUCTURES

Paris, France

21 Jul 2021 - 23 Jul 2021

ELLIS PhD and Postdoc Summit

Online via Zoom

12 Jul 2021 - 13 Jul 2021

Seminar "Research in videogames: use of deep learning for saliency estimation and cheating prevention"

30 Jun 2021 - 30 Jun 2021

Workshop on Artificial Scientific Discovery

29 Jun 2021 - 01 Jul 2021

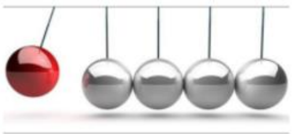

ELLIS Program "Interactive Learning and Interventional Representations" Workshop

21 Jun 2021 - 21 Jun 2021

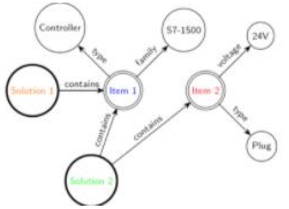

ELLIS Program "Semantic Symbolic, and Interpretable Machine Learning" Kick-off

07 Jun 2021 - 07 Jun 2021

ICRA 2021-Workshop on Semantic Representations for Robotics through Continuous Interaction and Incremental Learning

31 May 2021 - 31 May 2021

ELLIS-ESA workshop on quantum computing for huge data analysis, simulation and potential applications to Earth observation

27 May 2021 - 27 May 2021

ELLIS Human-Centric Machine Learning Workshop

10 May 2021 - 10 May 2021

French-German Machine Learning Symposium

10 May 2021 - 11 May 2021

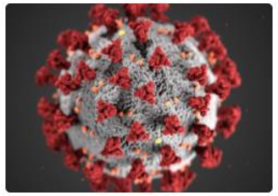

INTERNATIONAL VIRTUAL COVID-DATA CHALLENGE

28 Apr 2021 - 29 Apr 2021

PSL Intensive Weeks

Online

29 Mar 2021 - 02 Apr 2021

Foundations of Algorithmic Fairness

12 Mar 2021 - 16 Mar 2021

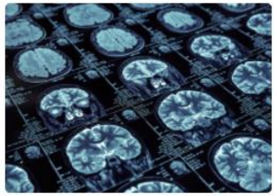

ELLIS Symposium Workshop on Geometric Deep Learning for Medical Imaging

02 Mar 2021 - 02 Mar 2021

Open Challenges and Future Directions of NLP

24 Feb 2021 - 25 Feb 2021

ELLIS Health Workshop: Explainable Machine Learning & Biological Mechanisms

Online

16 Dec 2020

Bi-weekly seminars on Automatic Machine Learning

03 Dec 2020

"About Time" - Seminar by Arnold Smeulders

30 Sep 2020

ELLIS Units: Official Launch

15 Sep 2020 - 15 Sep 2020

Workshop on Self-Supervised Robot Learning

13 Jul 2020

AI4Science Kickoff Workshop

08 Jul 2020

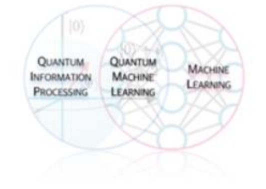

Online Workshop on Quantum and Physics based machine learning

No one size fits all: Artificial Intelligence as Key Technology for Personalised Medicine

Virtual Event

Global perspectives on responsible AI - An international and

CIFAR/ELLIS Discussion on Contact Tracing

23 Jun 2020

ELISE has issued its first open call for SMEs and Start-ups

Businesses that develop AI-based solutions and applications can apply for funding of up to 60,000 euros

ellis 29 April 2021 Announcement

ELISE open call funding SMEs Start-ups AI machine learning

The European Network of AI Excellence Centres (ELISE) will select 16 SMEs and start-ups in one of the ELISE/ELLIS focus areas.

The **European Network of AI Excellence Centres** (ELISE) is a network of artificial intelligence research hubs where the best European researchers in machine learning and AI work together to attract talent, foster research through collaboration, and inspire and be inspired by industry and society.

In its first open call, ELISE will select **16 SMEs and start-ups that develop AI services or applications.** The companies selected will take part in a six-month program and receive up to 60,000 euros in funding. While focus areas are based on ELISE/**ELLIS** research programs, proposals for projects that use machine learning to address major societal and economic challenges will also be considered. The deadline for applications is July 1, 2021.

More details can be found **here**. Join the live webinar on June 16 (12 PM CEST) for a Q&A; you can register **here**.

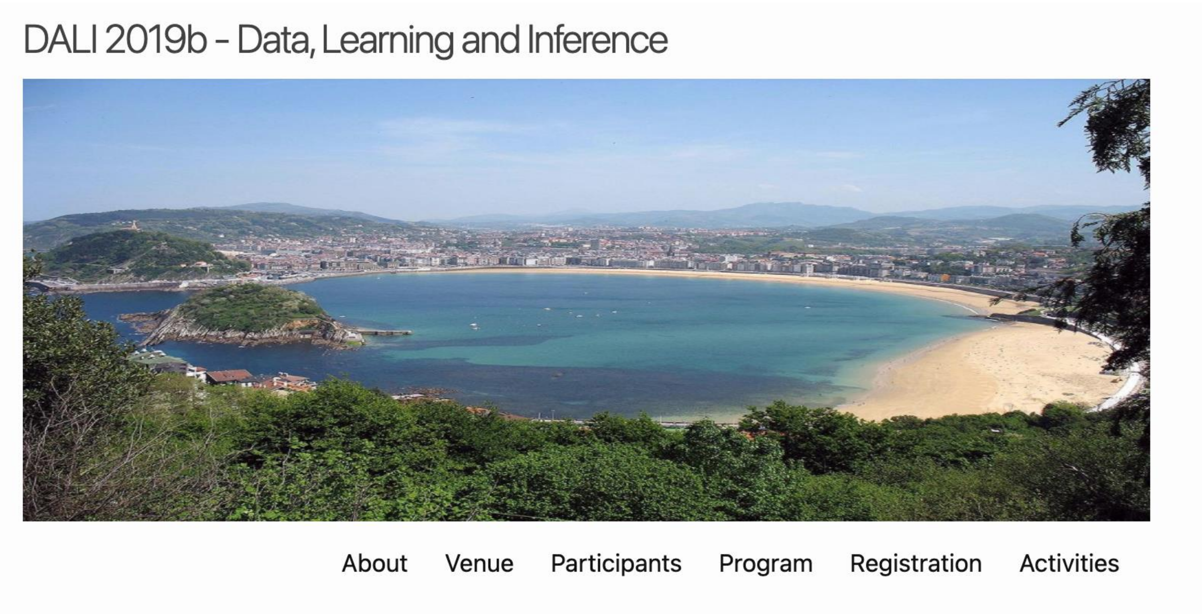
DALI 2019b - Data, Learning and Inference
About
Venue
Participants
Program
Registration
Activities


QPhML2020
Program
Organizers
Ellis QPhML 2020
Milan - Tübingen - Nijmegen, July 6-8
Ellis QPhML 2020
Registration is now open! Please register [here] with your email to attend the workshop.
The workshop will be held online. Details on how to connect will be announced soon.
e l l i s
European Laboratory for Learning and Intelligent Systems


EllisESA2021
Aims
Program
Organizers
e l l i s
ELLIS-ESA Workshop 2021
Quantum Algorithms and Machine Learning
for huge Data Analysis, Simulation
and potential Earth Observation Applications
esa
Ellis-ESA Workshop 2021
The workshop was held online. The videos of the talks can be found here
e l l i s
European Laboratory for Learning and Intelligent Systems
The Ellis program Quantum and Physics based machine learning (QPhML) is part of the recent European initiative called ELLIS (European Laboratory for Learning and Intelligent Systems) to stimulate research on machine learning by building networks of top research groups in Europe.
Learn more: [Ellis initiative, Ellis Fellows QPhML Program]
esa
Φ-lab

Ellis Fellows Program Quantum and Physics based Machine Learning (QPhML)

Ellis Fellows

Bert Kappen (program director) Department of Biophysics, Radboud University (Nijmegen). [web]

Riccardo Zecchina (program director) Department of Decision Sciences, Bocconi University (Milan). [web]

Miguel Angel Delgado Department of Theoretical Physics, Universidad Complutense (Madrid). [web]

David Gross Institute for Theoretical Physics, University of Cologne (Cologne). [web]

Florian Marquardt Institute for Theoretical physics, Max Planck Institute (Erlangen). [web]

Matthias Rupp Department of Theory, Fritz-Haber-Institut of the Max Planck Society (Berlin). [web]

Gabor Csanyi Department of Engineering, University of Cambridge (Cambridge). [web]

Florent Krzakala Department of Physics, Ecole Normal Superieur (Paris). [web]

Giulio Biroli Department of Physics, Ecole Normal Superieur (Paris). [web]

Lenka Zdeborova Institute for theoretical physics, University Paris-Saclay (Paris). [web]

Jens Eisert Dahlem Center for Complex Quantum Systems Free University Berlin (Berlin). [web]

Giuseppe Santoro SISSA (Trieste). [web]

Remi Monasson Department of Physics, Ecole Normal Superieur (Paris). [web]

Carlo Baldassi Department of Decision Sciences, Bocconi University (Milan). [web]

Vedran Dunjko Leiden Institute for Advanced Computer Science, University Leiden (Leiden). [web]

Giuseppe Carleo Center for Computational Quantum Physics, Flatiron Institute (New York). [web]

Marc Mezard Department of Physics, Ecole Normal Superieur (Paris). [web]

Nicolas Regnault Department of Physics, Ecole Normal Superieur (Paris). [web]

Jorge Kurchan Department of Physics, Ecole Normal Superieur (Paris). [web]

Matthias Troyer ETH Zurich and Microsoft Research. [web]

Manfred Opper TU (Berlin). [web]

Hans Briegel Institut für Theoretische Physik, University of Innsbrück (Innsbrück). [web]

Aram Harrow MIT (Boston). [web]

Valentina Ros Department of Physics, Ecole Normal Superieur (Paris). [web]

Andrea Rocchetto Department of Computer Science, University of Texas at Austin, (Austin). [web]

Frank Noe AI4Science, Freie Universitaet (Berlin). [web]

Quantum enhanced machine learning

> Quantum devices are nearing the noisy intermediate scale quantum (NISQ) era, corresponding to machines with 50 to 100 qubits and capable of executing circuits with depths on the order of thousands of elementary two qubit operations. NISQ devices may provide computational advantages over classical supercomputers for various machine learning problems, which includes sampling from hard-to-simulate probability distributions for Bayesian methods and the Quantum Boltzmann Machine and linear algebra problems (for instance for kernel methods or deep learning). It is hoped that the application of NISQ technology to machine learning may be one of the first instances exhibiting genuine quantum advantages.

Statistical physics approach to machine learning

> Noise plays a fundamental role for learning in large neural networks. Rather than designing reliable bits and use software to generate random numbers, an appealing alternative is to design hardware that is noisy by design. Such devices would be much more energy efficient. Methods from non-equilibrium statistical physics are well suited to improve our understanding of stochastic systems. An example is the use of physically coupled replicas that have been shown to be very effective for hard combinatoric or strongly non-linear learning problems. In addition, the observation that physical replicas resemble Trotterized quantum systems provides a promising new research direction for the design of stochastic or quantum learning algorithms. Another link between quantum and stochastic systems is the observation that sign free quantum systems can be mapped onto classical stochastic diffusion problems.

Using machine learning for quantum physics

> The challenge of quantum many-body physics is to efficiently describe and control exponential numbers of parameters of quantum systems. Better characterization of such systems will lead to the understanding of quantum materials such as high-temperature superconductors or topological insulators. Enhanced control of immense parameter spaces will improve the understanding and design of quantum devices, enabling quantum computers and networks. For this problem, machine learning offers a new option.

Quantum based ML

- **Quantum learning theory** (Aram Harrow, Vedran Dunjko, Andrea Rochetto, Jens Eisert)
- **Tensor networks for ML** (Jens Eisert, Hans Briegel)
- **Variational/Parametrized circuits** (Marcello Benedetti, Hans Briegel, Jens Eisert, Christian Gogolin, Aram Harrow)
- **Quantum Boltzmann Machine** (Bert Kappen, Leonard Wossnig)
- **Autonomous learning in classical/quantum systems** (Bert Kappen, Floriant Marquardt)
- **Quantum applied machine learning** (Giuseppe Carleo, Gabor Csanyi, Christian Gogolin, Florian Marquardt)
- **Quantum Chemistry** (Mathias Rupp, von Lilienfeld, Leonard Wossnig, Frank Noe)
- **Quantum RL** (Hans Briegel, Vedran Dunjko)

Quantum Learning Theory

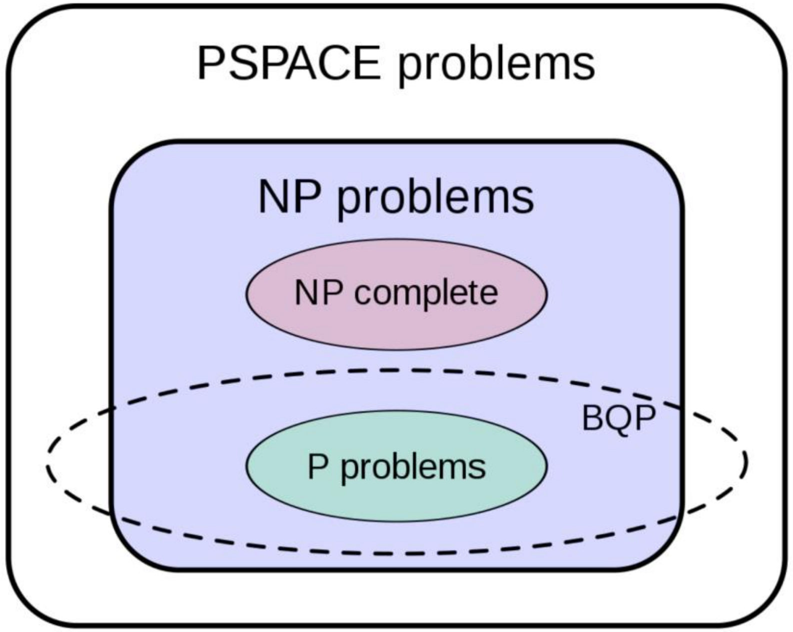

Image credits: wikipedia.org

P: solvable in poly time
NP: solution verifiable in poly time
Pspace: solvable with poly memory
BQP: Bounded Quantum Polynomial

It is conjectured that BQP solves hard problems outside of P, specifically, problems in NP. Examples are

- Integer factorization (Shor's algorithm) $\mathcal{O}\left(e^{N^{\frac{1}{3}}}\right) \rightarrow \mathcal{O}\left(N^2\right)$
- Solving sparse linear system (HHL) $\mathcal{O}\left(N\right) \rightarrow \mathcal{O}\left(\log N\right)$

Aram Harrow, Vedran Dunjko, Andrea Rochetto, Jens Eisert)

Quantum applied machine learning

Solving the Quantum Many-Body Problem with Artificial Neural Networks

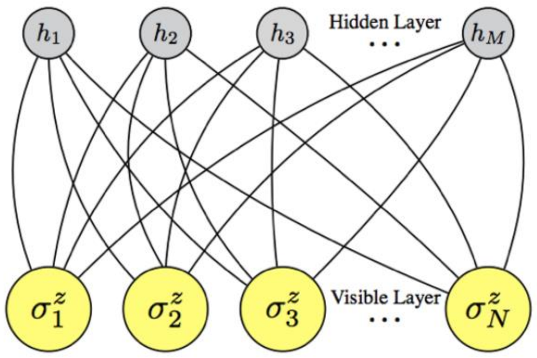

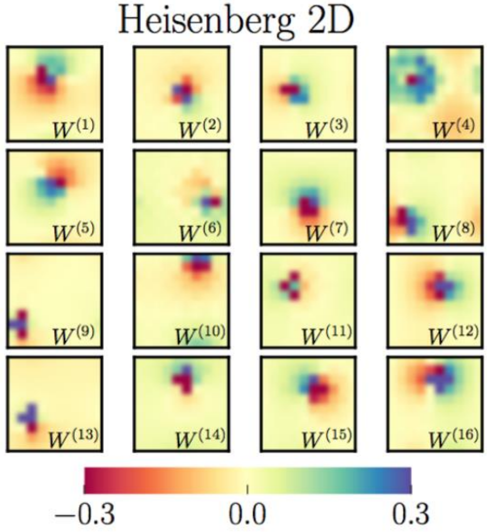

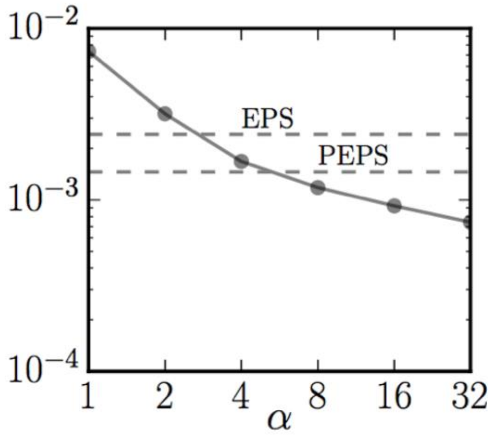

Carleo, Troyer, Science 2017

Variational/parametrized circuits

$$\Phi : \vec{x} \in \Omega \to |\Phi(\vec{x})\rangle\langle\Phi(\vec{x})|$$

Supervised learning with quantum enhanced feature spaces

Vojtech Havlicek[1,*] Antonio D. Córcoles[1], Kristan Temme[1], Aram W. Harrow[2],
Abhinav Kandala[1], Jerry M. Chow[1], and Jay M. Gambetta[1]
[1] *IBM T.J. Watson Research Center, Yorktown Heights, NY 10598, USA and*
[2] *Center for Theoretical Physics, Massachusetts Institute of Technology, USA*
(Dated: June 7, 2018)

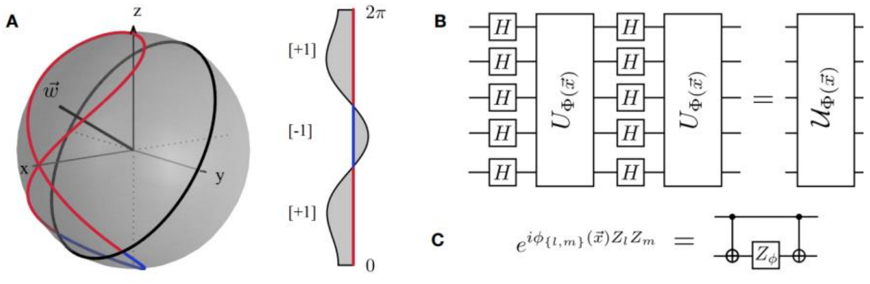

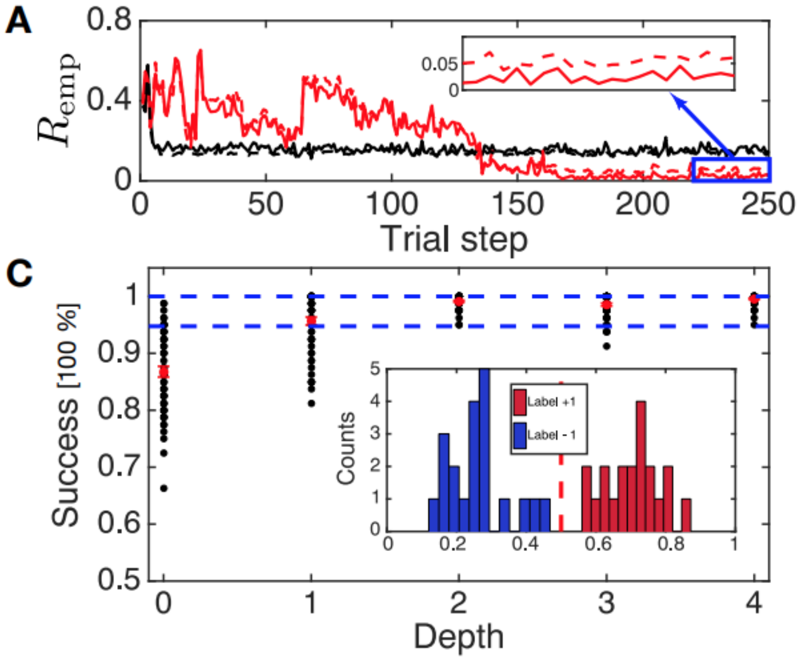

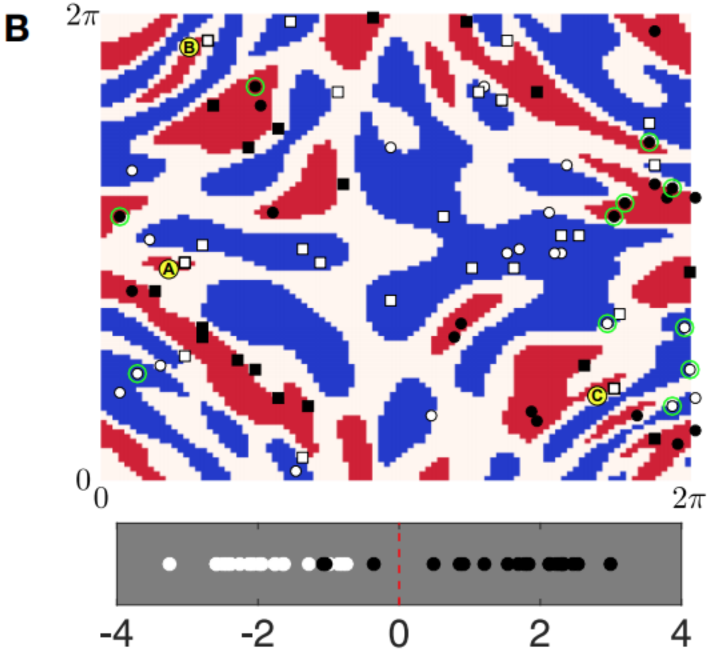

Quantum Boltzmann Machine

"Can we design a Hamiltonian such that its ground state represents a given (data) distribution"

This is a learning problem.

The quantum state is a density matrix

$$\rho = \frac{1}{Z} e^{H} \qquad Z = \mathrm{Tr}\left(e^{H}\right) \qquad H = \sum_r H_r w_r$$

with w_r free parameters.

The target is a rank-1 density matrix, computed from the data distribution q:

$$\eta = |\phi\rangle \langle\phi| \qquad \langle s|\phi\rangle = \sqrt{q(s)}$$

Learning is to find w_r that minimize the relative entropy $S(\eta, \rho)$ by gradient descent:

$$\Delta w_r = -\epsilon \frac{\partial S}{\partial w_r}$$

Modeling neural network activity

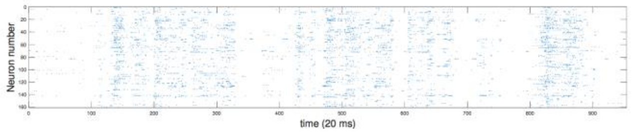

(a) One repeat of neural activity of 160 salamander retinal ganglion cells [Schneidman et al., 2006].

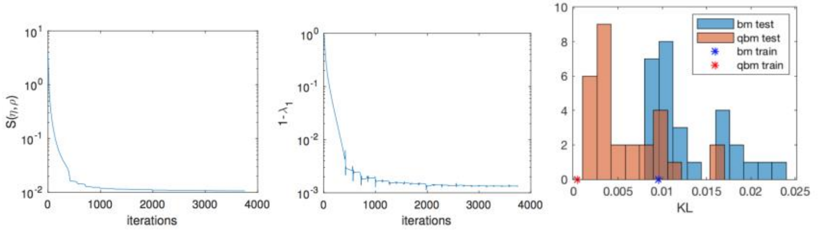

(b) QBM Training performance. (c) QBM Rank-1-ness. (d) BM and QBM test performance on 28 independent test sets.

Parity problem

Data are generated such that even parity states have $q(s) > 0$ and odd states have $q(s) = 0$.

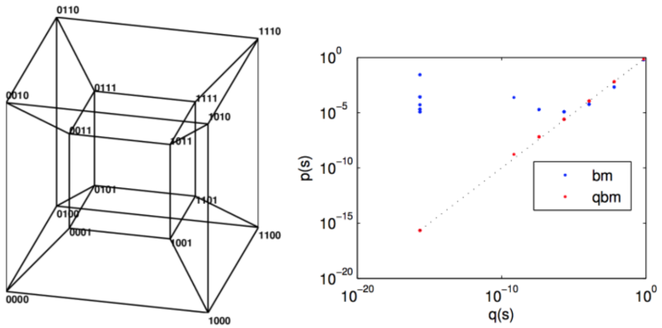

QBM learns a rank 1 solution: $KL(q|p_{\mathrm{qbm}}) = 1.31 \times 10^{-5}$
BM cannot learn this problem: $KL(q|p_{\mathrm{bm}}) = 0.451$

Kappen 2020

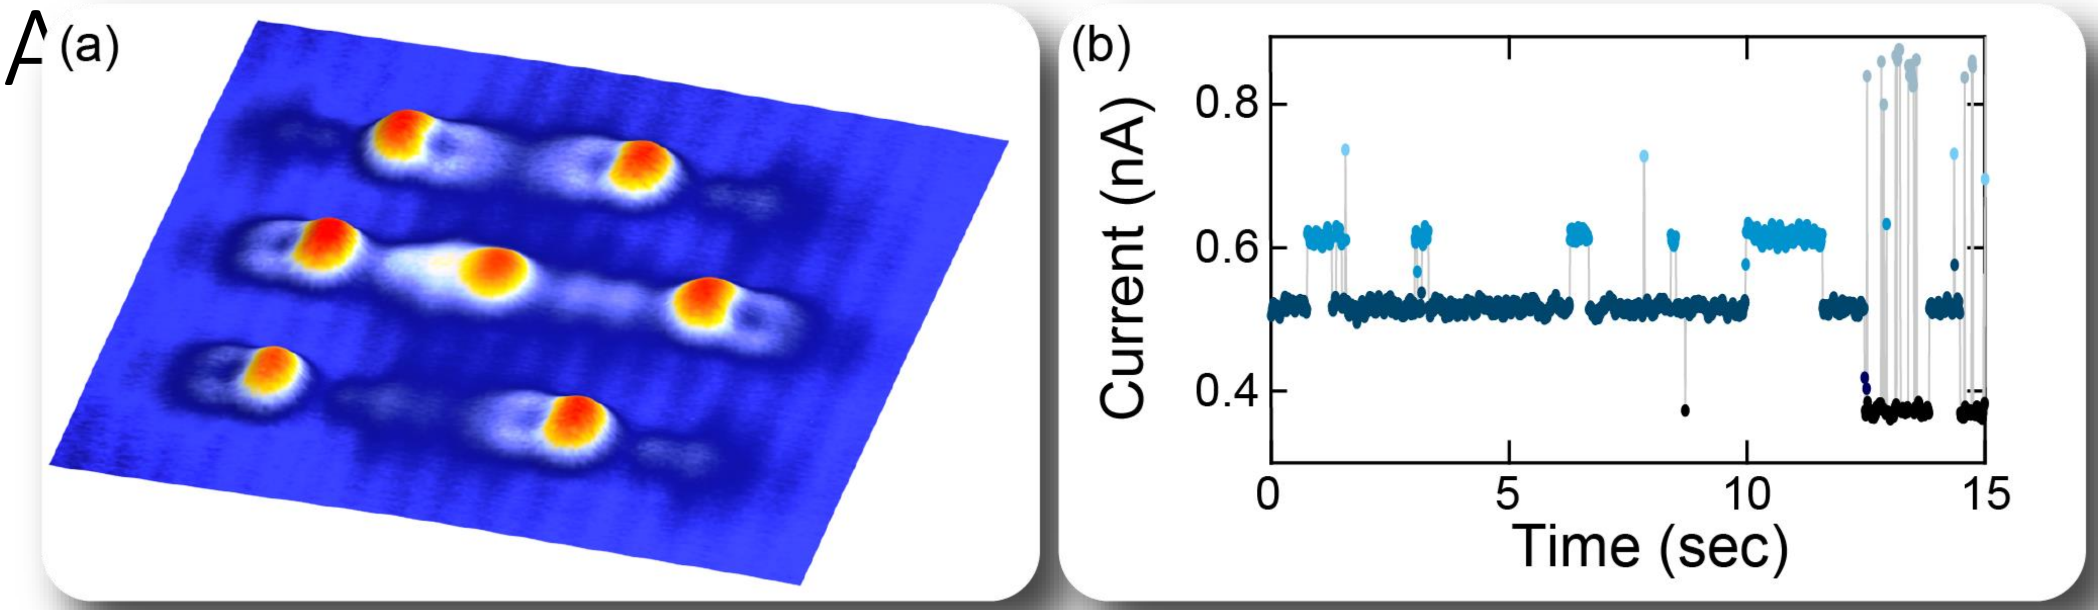

- Surface dopant network, based on orbital memory – may be high temperature compatible (currently 4K)
- Large separation of time scales, leads to integrated neurons/synapses
- Material exhibits self-adaption
- Spin-based properties and response to external fields, unexplored.
- Radboud consortium: Khajetoorians, Kappen, Katsnelson
    - Links to Twente: van der Wiel

Physics based ML

- Statistical physics of learning (Riccardo Zecchina, Carlo Baldassi, Marc Mezard, Florent Krzakala, Giulio Biroli, Lenka Zdeborova, Remi Monasson)
  - Random satisfiability
  - Error correction
  - Compressed sensing
  - Entropy based learning for binary perceptron

Random satisfiability [Mezard et al., 2002]

The K-satisfiability problem (K-sat) asks whether one can satisfy simultaneously a set of M constraints between N Boolean variables $x_i = 0, 1$, where each constraint is a clause built as the logical OR involving K variables. An instance of 3-sat is

$$(\neg x_1 \vee x_2 \vee \neg x_3) \wedge (x_1 \vee \neg x_4 \vee x_5) \wedge \ldots$$

3-sat is at the core of combinatorial optimization theory. An efficient algorithm for solving 3-sat would immediately lead to other algorithms for efficiently solving thousands of different NP hard combinatorial problems [Garey and Johnson, 1979].

Optimisation problems

Assignment (“easy”, in P)

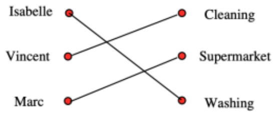

Travelling salesman (“hard”, NPC)

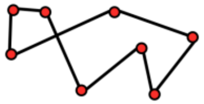

Hamiltonian path (“hard”, NPC)

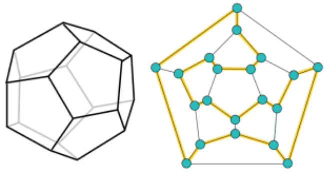

Random satisfiability [Mezard et al., 2002]

Define $\alpha = M/N$. K-sat has many solutions when α is small (SAT phase) and no solutions when α is large (UNSAT phase). Random 3-sat can be analyzed with the cavity method.

Fig. 2. The phase diagram of the random 3sat problem. Plotted is e_0, the number of violated clauses per variable (red), versus the control parameter α, which is the number of clauses per variable. The SAT-UNSAT transition occurs at $\alpha = \alpha_c \sim 4.256$. The green line is e_{th}, the threshold energy per variable, where local algorithms get trapped. The blue line is the complexity Σ of satisfiable states, equal to $1/N$ times the logarithm of their number.

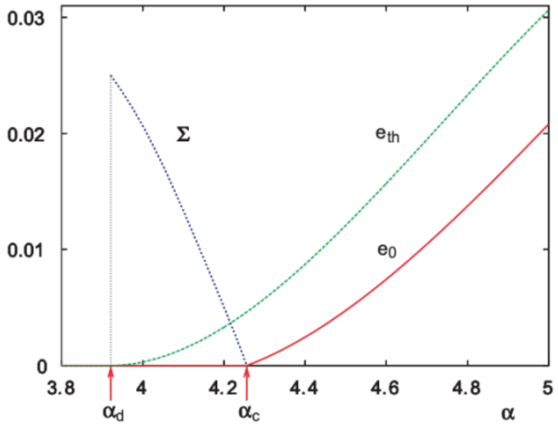

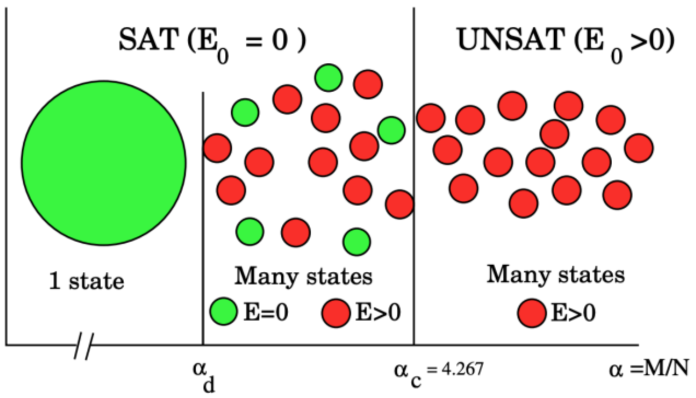

(Red) Lowest energy $e_0 = E_0/N$ versus α. When $e_0 > 0$ the problem is UNSAT. (Green). Lowest energy $e_{\text{th}} = E_{\text{th}}/N$ that can be obtained with local algorithms. (Blue) Number of states $e^{N\Sigma}$ with $E = 0$.

When $\alpha_d < \alpha < \alpha_c$, the clauses cause frustration, similar to the SK spin glass: there are local minima either $E = 0$ or $E > 0$ and it is hard to find a $E = 0$ solution using local methods (such as iterative improvement or simulated annealing).

Random satisfiability [Mezard et al., 2002]

The simplest algorithm to solve the 3-sat problem is to use the max product (or equivalently max sum) on the factor graph. Because of the zero-one nature of the energy, each message $m_{i \to a}(x_i)$ or $m_{a \to i}(x_i)$ has only two values ± 1.

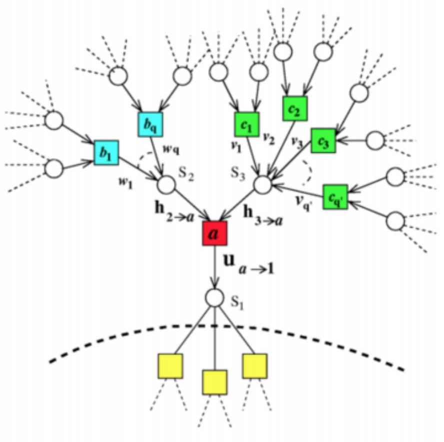

The survey propagation algorithm generalizes this to three values $(-1, 0, 1)$. Furthermore, a message is a distribution over these three values.

The SID (survey inspired decimation) algorithm is an iterative prodedure where in each iteration a subset of variables are clamped to ± 1 based on the survey progapation result. In the hard regime ($\alpha = 4.2$) SID confirms the solution on existing benchmarks for $N = 2000$. SID obtains solutions up to instances of size $N = 100.000$ where no other method can be appied. The complexity of SID is quadratic in N.

See [Mezard and Zecchina, 2002] for details.

Compressed sensing

Given an unknown signal which is a N-dimensional vector $x = (x_1, \ldots, x_N)$, we make M measurements $y_i = \sum_{j=1}^{N} F_{ij} x_j$. For instance, measurements of Fourier modes or wavelet coefficients. The observer knows the $M \times N$ matrix F and the measurements y. His aim is to reconstruct x.

When $M = N$ the solution is obtained by matrix inversion: $x = F^{-1} y$.

When $M < N$ the problem is underdetermined and there are many solutions.

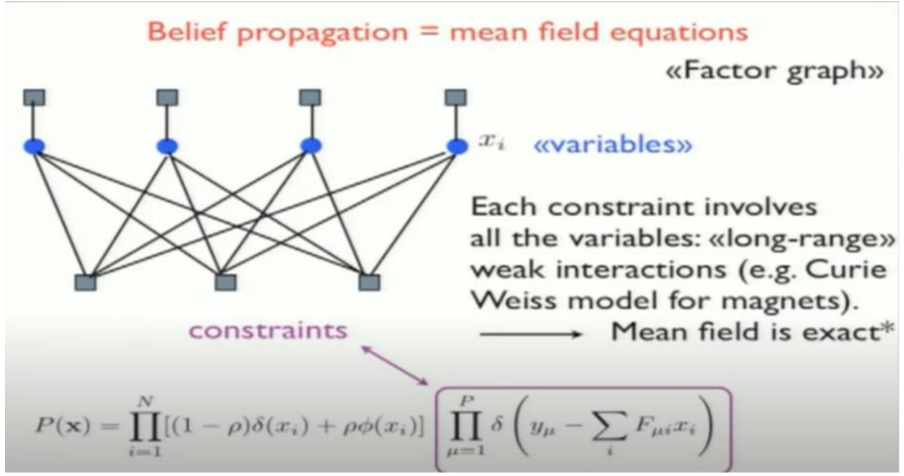

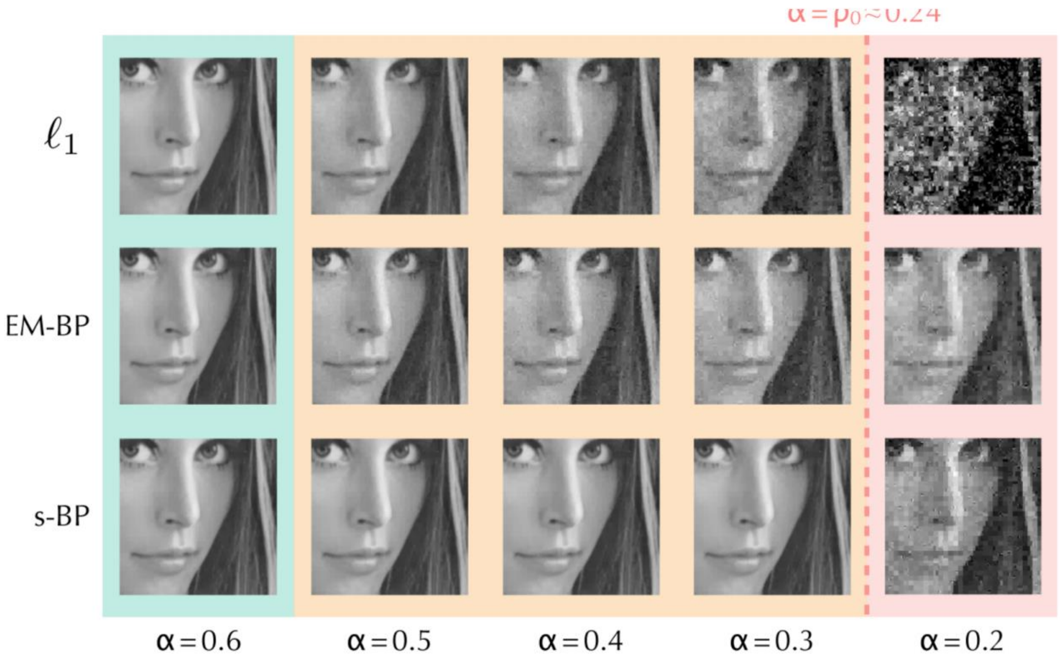

[Krzakala et al., 2012]

Binary perceptron

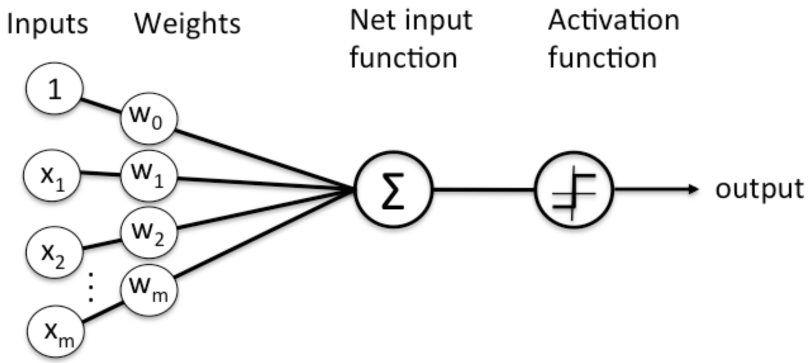

Schematic of Rosenblatt's perceptron.

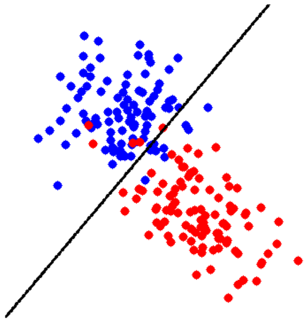

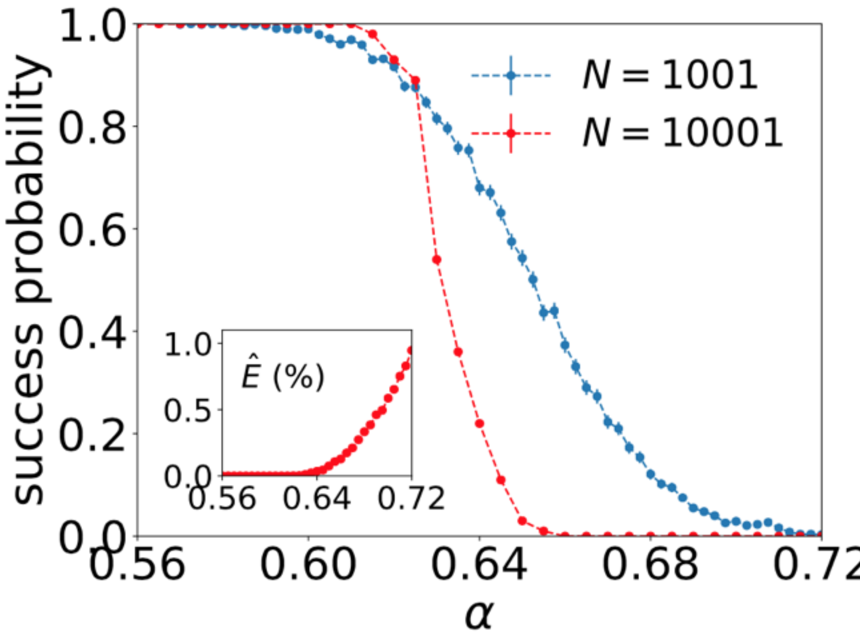

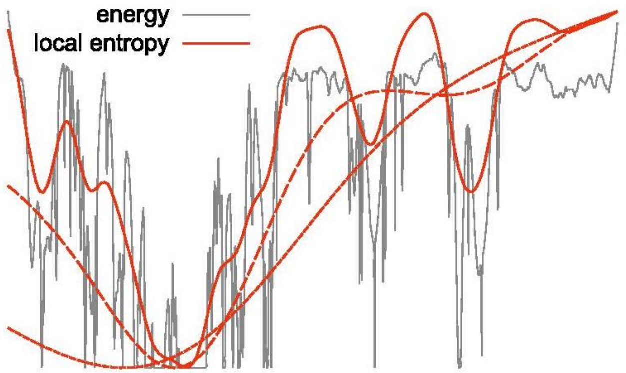

Continuous Perceptron learning is easy and can learn 2N patterns.
Binary perceptron learning is NP hard and can learn 0.83N patterns

Optimization is hard, is improved through local entropy measure, which
Is mapped on a replicated system

MNIST

- On the MNIST benchmark problem, using a network with three hidden layers we achieved ~ 1.7% test error, a very good result for a network with binary weights and activations and with no convolutional layers.

Future aim QPhML

- Integrate physics and machine learning (quantum and stat phys)
  - Organize annual meeting
  - Special topics meetings (ESA, CERN?)
  - Special issues