Scientific Program

July 4

9:00-9:15	Registration
9:15-9:30	Opening
9:30-10:30	**Eun-Ah Kim (Cornell)** *Learning Quantum Emergence with AI*
10:30-11:00	Break
11:00-12:00	**Giuseppe Carleo (Flatiron Institute)** *Neural-Network Quantum States:* *from Condensed Matter to Quantum Computing*
2:00-3:00	**Kieron Burke (UC Irvine)** *Creating New Density Functionals with Machine-learning*
3:00-4:00	**Lexing Ying (Stanford)** *Solving PDEs with Deep Learning*
4:00-4:30	Break
4:30-5:30	**Gábor Csányi (Cambridge)** *A New Dawn of Interatomic Potentials*

July 5

9:30-10:30	**Hans J. Briegel (Innsbruck)** *Machine learning for the autonomous design* *of quantum physics experiments*
10:30-11:00	Break
11:00-12:00	**Evert van Nieuwenburg (Caltech)** *How Confused is My Network?*
2:00-3:00	**Matthias Rupp (FHI Berlin)** *Machine Learning for Interpolation of Electronic Structure Calculations*
3:00-4:00	**José Miguel Hernández-Lobato (Cambridge)** *Advances in Machine Learning for Molecules*
4:00-4:30	Break
4:30-5:30	**Lei Wang (Institute of Physics, CAS)** *Neural Network Renormalization Group*
5:30-6:30	Poster Session

July 6

9:30-10:30	**Fakher F. Assaad (Wuerzburg)** *Monte Carlo Simulations of Quantum Matter*
10:30-11:00	Break
11:00-12:00	**Huitao Shen (MIT)** *Boosting Quantum Monte Carlo Simulations with Machine Learning*
2:00-3:00	**Jim Halverson (Northeastern)** *Reinforcement Learning and the String Landscape*
3:00-4:00	**Daniel Roberts (Facebook AI Research)** *Why is AI hard and Physics Simple?*
4:00-4:30	Break
4:30-5:30	**Yizhuang You (UCSD)** *Machine Learning Holography*
5:30-6:30	Discussion and Concluding

Titles and Abstract of the Talks

July 04, Wednesday

Learning Quantum Emergence with AI

Eun-Ah Kim, Cornell

I will discuss 1) using machine learning to learn from big data on quantum materials (STM), 2) interpreting what machine learned in obtaining a phase diagram from simulation of topological quantum phase transition.

Neural-Network Quantum States:

from Condensed Matter to Quantum Computing

Giuseppe Carleo, Flatiron Institute

Machine-learning-based approaches, routinely adopted in cutting-edge industrial applications, are being increasingly adopted to study fundamental problems in science. Very recently, their effectiveness has been demonstrated also for many-body physics.

In this seminar I will present recent applications to the quantum realm. First, I will discuss how a systematic machine learning of the many-body wave-function can be realized. This goal has been achieved in [1], introducing a variational representation of quantum states based on artificial neural networks. This representation can be used to study both ground-state and unitary dynamics, with controlled accuracy. I will then show how a similar representation can be used for applications directly relevant to ultra-cold atoms and quantum computing. In this context, I will discuss both Quantum State

Tomography of highly-entangled states [2], and a novel approach for the classical simulation of large quantum circuits [3].

[1] Carleo, and Troyer — Science 355, 602 (2017).
[2] Torlai, Mazzola, Carrasquilla, Troyer, Melko, and Carleo — Nature Physics 14, 447-450 (2018).
[3] Jonsson, and Carleo — In preparation (2018)

Creating New Density Functionals with Machine-learning

Kieron Burke, UC Irvine

This talk is designed to be accessible to folks with a wide variety of backgrounds. In the first half, I will discuss the growing interest in data-enabled chemistry, including a special issue of the journal of chemical physics on the topic[2]. In the second part, I will briefly review density functional theory and why it is important to many branches of modern physical science[3]. And in the third half, I will show how, in collaboration with computer scientists at TU Berlin, we have used a specific type of machine-learning, called kernel ridge regression, to find more accurate and powerful approximate density functionals than any made by humans[1]. Lastly, I will describe work in collaboration with Steve White, using DMRG to train a machine-learned functional to do strongly correlated systems.

References:

1. By-passing the Kohn-Sham equations with machine learning Felix Brockherde, Leslie Vogt, Li Li, Mark E Tuckerman, Kieron Burke, Klaus-Robert Muller, Nature Communications 8, 872 (2017).

2. Special Topic in Data-enabled Chemistry, https://aip.scitation.org/toc/jcp/148/24

3. DFT in a nutshell Kieron Burke, Lucas O. Wagner, Int. J. Quant. Chem. 113, 96-101 (2013)

Solving PDEs with Deep Learning

Lexing Ying, Stanford

In this talk, I will discuss some recent work on using deep neutral networks in solving high dimensional PDE problems. Examples include homogenization, density functional theory, and molecular dynamics.

A New Dawn of Interatomic Potentials

Gábor Csányi, Cambridge

I will show our recent work on data driven interatomic potentials. The goal of this research programme is to construct analytic functions that accurately reproduce the Born-Oppenheimer potential energy surface of condensed phase materials. Much progress has been made by an increasing number of groups over the last few years, mostly by borrowing approaches and attitudes from the field of machine learning - even though the mathematical context is rather different. Accurate potentials have been published by us for carbon, silicon, boron, tungsten, iron, that cover a wide range of atomic environments, and for many other materials by other groups. These potentials are beginning to be used in materials science applications.

Recent References:

1) Bartók, AP and De, S and Poelking, C and Bernstein, N and Kermode, JR and Csányi, G and Ceriotti, M (2017) Machine learning unifies the modeling of materials and molecules. Sci Adv, 3. e1701816

2) Deringer, VL and Csányi, G (2017) Machine learning based interatomic potential for amorphous carbon. Physical Review B, 95. ISSN 2469-9950

3) Dragoni, D and Daff, TD and Csanyi, G and Marzari, N (2018) Achieving DFT accuracy with a machine-learning interatomic potential: Thermomechanics and defects in bcc ferromagnetic iron. Physical Review Materials, 2. ISSN 2475-9953

4) Deringer, VL and Pickard, CJ and Csányi, G (2018) Data-Driven Learning of Total and Local Energies in Elemental Boron. Phys Rev Lett, 120. 156001

5) Caro, MA and Deringer, VL and Koskinen, J and Laurila, T and Csányi, G (2018) Growth

Mechanism and Origin of High sp^{3} Content in Tetrahedral Amorphous Carbon. Phys Rev Lett, 120. 166101

July 05, Wednesday

Machine learning for the autonomous design of quantum physics experiments

Hans J. Briegel, Innsbruck

In this talk, I will discuss the use of artificial learning agents in quantum physics laboratories, as well as the use of quantum information in machine learning and artificial-agent design. I will focus on the model of projective simulation (PS) [1], which employs random-walk processes in the agent's memory for learning and decision-making. Projective simulation has been applied, e.g., in autonomous robotic playing [2] and in the design of quantum experiments [3]. The PS model can be naturally quantized, allowing for a quantum speed-up of the agent's decision process [4]. I will review some recent results of our research on (classical and) quantum-enhanced learning agents [5], including applications in quantum optical experiments and quantum foundations.

Literature:

1) H. J. Briegel and G. De las Cuevas (2012) Projective simulation for artificial intelligence, Scientific Reports 2, 400.
2) S. Hangl, E. Ugur, S. Szedmak, and J. Piater (2016) Robotic playing for hierarchical complex skill learning, Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference, pp. 2799 – 2804.
3) A. Melnikov, H. Poulsen Nautrup, M. Krenn, V. Dunjko, M. Tiersch, A. Zeilinger, and H. J. Briegel (2018) Active learning machine learns to create new quantum experiments, PNAS 115, 1221.
4) G. Paparo, V. Dunjko, A. Makmal, M. A. Martin-Delgado, and H. J. Briegel (2014) Quantum speed-up for active learning agents, Phys. Rev. X 4, 031002.
5) V. Dunjko, J. M. Taylor, and H. J. Briegel (2016) Quantum-enhanced machine learning, Phys. Rev. Lett. 117, 130501.

How Confused is My Network?

Evert van Nieuwenburg, Caltech

In this contribution I will discuss the 'learning by confusion' method, and its extension to 'discriminative cooperative networks'. Both of these methods rely on the ability to order physics data along one or multiple tuning parameters; a feature not always possible in non-physics data-sets. In the second part, I will consider the use of recurrent neural networks to construct a phase diagram from dynamics of observables.

Machine Learning for Interpolation of Electronic Structure Calculations

Matthias Rupp, FHI Berlin

Systematic computational study, discovery and design of novel molecules and materials requires accurate simulations on the atomic scale. While numerical approximations to the electronic structure problem enable this in principle, their applicability is severely limited by their high computational cost. In high-throughput settings, machine learning can reduce these costs significantly by interpolating between reference calculations. Effectively, the problem of solving a complex equation such as the electronic Schrödinger equation for many related inputs is mapped onto a nonlinear statistical regression problem. I will provide an introduction to kernel-based machine learning approaches for the accurate and rapid interpolation of electronic structure calculations. For such approaches, a numerical representation of atomistic systems that supports interpolation is crucial. Using our recently introduced many-body tensor representation, I will present empirical evidence for accurate predictions of ab initio formation enthalpies on diverse datasets of molecules and crystal structures.

Advances in Machine Learning for Molecules

José Miguel Hernández-Lobato, Cambridge

In this talk, I will describe two applications of machine learning to molecule data. First, I will focus on the problem of efficiently searching chemical space for new molecules with optimal properties. I will describe how to use recent advances in deep generative models to obtain continuous representations of molecules which allow us to automatically generate novel chemical structures by performing simple operations in a latent space. These methods can then be connected with Bayesian optimization techniques to accelerate the search for new molecules with optimal properties. In the second part of the talk, I will focus on the problem of modeling chemical reactions by predicting electron paths. Chemical reactions can be described as the stepwise redistribution of electrons in molecules. As such, reactions are often depicted using "arrow-pushing" diagrams which show this movement as a sequence of arrows. I will describe an electron path prediction model to learn these sequences directly from data and show that the model recovers a basic knowledge of chemistry without being explicitly trained to do so.

Neural Network Renormalization Group

Lei Wang, Institute of Physics, CAS

Abstract: I will present a variational renormalization group (RG) approach using a deep generative model based on normalizing flows. The model performs a hierarchical of change-of-variables transformations from the physical space to a latent space with reduced mutual information. Conversely, it directly generates statistically independent physical configurations as a form of inverse RG flow. The generative model has an exact and tractable likelihood, which allows unbiased training and direct access to the renormalized energy function of the latent variables. To train the neural network, we employ the probability density distillation of the bare energy function, where the training loss provides a variational upper bound of the physical free energy. We demonstrate practical usage of the approach by identifying mutually independent collective variables of the Ising model and performing accelerated hybrid Monte Carlo sampling in the latent space. I will comment on the connection of the present approach to DeepMind's WaveNet, wavelet formulation of RG, and the modern pursuit of information preserving RG.

Reference: arXiv:1802.02840

July 06, Wednesday

Monte Carlo Simulations of Quantum Matter

Fakher F. Assaad, Wuerzburg

The auxiliary field quantum Monte Carlo approach is a method of choice to unravel emergent collective phenomena in a set of many electron systems ranging from the solid state to particle physics. In fact, recent progress in our understanding of the class of models that can be simulated without encountering the negative sign problem renders this approach very powerful. There are however many challenges, such as automatic generation of phase diagrams and sampling strategies. In this talk, I will review the approach and place emphasis on how concepts of machine learning can enhance the efficiency of the algorithm.

Boosting Quantum Monte Carlo Simulations with Machine Learning

Huitao Shen, MIT

Quantum Monte Carlo is a large family of powerful numerical tools to understand the quantum many-body systems at zero temperature and finite temperatures. It plays an important role in condensed matter physics and quantum chemistry. In this talk, we show how to incorporate various machine learning techniques to improve the performance of these quantum Monte Carlo algorithms.

Reinforcement Learning and the String Landscape

Jim Halverson, Northeastern

String compactifications to four dimensions give rise to a large ensemble of metastable ground states with differing physical excitations that are determined in part by the topology and geometry of extra dimensions. This is the landscape of string theory. Recently, numerous techniques in data science have

been used to study physics that arises in the landscape. In this talk I will focus on the application of reinforcement learning to a number of problems, including to the problem of finding small cosmological constants.

Why is AI hard and Physics Simple?

Daniel Roberts, Facebook AI Research

We discuss why AI is hard and why physics is simple. In particular, we discuss how physical intuition and the approach of high-energy theoretical physics can be brought to bear on the field of artificial intelligence, and specifically machine learning. To that end, we suggest that the underlying project of machine learning generalizes the project of physics.

To make this point concrete, we focus on the dynamics of gradient descent. First, we interpret gradient descent in its continuum limit by giving an action formulation. Next, we discuss a toy model of gradient descent in high-dimensional spaces, and we argue that the gradient dynamically converges to a very small subspace after a short period of training. A simple argument then implies that gradient descent happens mostly in this subspace. We show experimental results indicating that this effect occurs realistically in a wide variety of large-scale deep learning scenarios, and we comment on possible implications for optimization and learning.

Machine Learning Holography

Yi-Zhuang You, UCSD

There is a profound relation between renormalization group (RG), holographic duality and deep learning. An information preserving RG is equivalent to a generative model, and the correspondence between the features in the dataset and the generative model that captures these features can be considered as a holographic duality. As a concrete example, we focus on the entanglement features of a quantum many-body state, and show that how holographic geometry can emerge as a result of deep learning the entanglement features.
