Physics & AI Workshop

Explore how artificial intelligence can be used in the physical sciences.

Conference Booklet

May 6th – 7th 2019

ORS Object Research Systems

TANDEMLAUNCH

RQMP

IVADO

Calcul Québec

McGill HPC

Institut Quantique Université de Sherbrooke

McGill Department of Physics

CAMBAM Centre for Applied Mathematics in Bioscience and Medicine

SCHEDULE

Monday

		Room
08:00 - 08:45	Registration	Foyer
08:45 - 09:00	Welcome	Ballroom A
09:00 - 09:45	Keynote - Roger Melko	Ballroom A
09:45 - 10:15	Coffee	Foyer
10:15 - 11:15	Short talks- TandemLaunch, Hilke, Fort, Ivanov, Granato	Ballroom A
11:15 - 12:00	Keynote - Edward Ott	Ballroom A
12:00 - 13:00	Lunch	Foyer
13:00 - 14:00	Short talks-ORS, Ouqamra, Kolosova, Lavoie, Beeler	Ballroom A
14:00 - 16:00	Workshop	Beginner: Ballroom A - Intermediate: Salle Du Parc
16:00 - 18:00	Reception/poster exhibit	Ballroom B

Tuesday

08:00 - 09:00	Breakfast	Foyer
09:00 - 09:45	Keynote - Pierre Bellec	Ballroom A
09:45 -10:15	Short talks-Martz-Oberlander (WiP), Whittaker	Ballroom A
10:15 - 10:45	Coffee	Foyer
10:45 - 11:15	Short talks- Bottrel, Mukkavilli	Ballroom A
11:15 - 12:00	Keynote - Guillaume Lamoureux	Ballroom A
12:00 - 13:15	Lunch	Foyer
13:15 - 14:00	Invited talk - Hong Guo	Ballroom A
14:00 - 16:00	Workshop	Beginner: Ballroom A - Intermediate: Salle Du Parc
16:00 - 18:00	Reception/poster exhibit	Ballroom B

Slack

We invite you to join the slack channel. Announcements and files for the workshop will be shared there. Invitation link: `https://bit.ly/2GRkAiK`

Floorplan

New Residence Hall, 3625 Av du Parc, Montral, QC H2X 3P8. The talks will take place in the basement ballroom A.

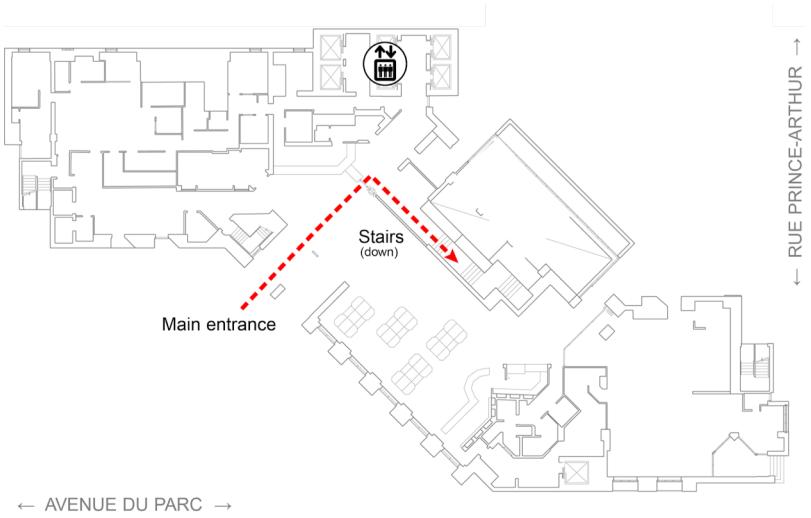

Figure 1: Ground floor

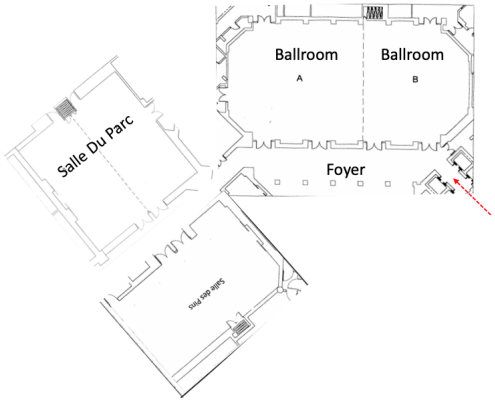

Figure 2: Basement

Welcome Address

In recent years, artificial intelligence has made remarkable advancements, impacting many industrial sectors dependent on complex decision-making and optimization. Physics-leaning disciplines also face hard inference problems in complex systems: climate prediction, density matrix estimation for many-body quantum systems, material phase detection, protein-fold quality prediction, parametrization of effective models of high-dimensional neural activity, energy landscapes of transcription factor-binding, etc. Methods using artificial intelligence have in fact already advanced progress on such problems. So, the question is not whether, but how AI serves as a powerful tool for data analysis in academic research, and physics-leaning disciplines in particular.

With the Physics AI Workshop, we aim to (1) stimulate the research community at the intersection of physics and machine learning in Montreal, and (2) provide researchers hoping to learn more about artificial intelligence with knowledge of state-of-the-art applications of machine learning in the physical sciences as well as hands-on training using standard tool sets.

Contents

Slack . . . iii

Floorplan . . . iii

Welcome Address **iv**

Workshop Outlook **1**

Beginner's Workshop (***Evgeny Naumov***) . . . 1

Intermediate Workshop (***Jean-Philippe Reid***) . . . 1

Keynote 1 **2**

Generative Models as AI-Assistants for NISQ Hardware (***Roger Melko***) . . . 2

Short talk session 1 brought to you by Tandem Launch: **3**

AI + Physics : United to protect your privacy online (***Alexandre Desilets-Benoit***) . . . 3

(Quantum) Machine Learning Applied to Disordered Systems (*D. Bernardi, A. Cheng, D. Ittah, D. Lopez, and A. Vuong, and* ***M. Hilke***) . . . 3

Adaptive Quantum State Tomography with Neural Networks (***Stanislav Fort****, Yihui Quek, Hui Khoon Ng*) . . . 4

Efficient Method for Solving the Time-Dependent Schrodinger Equation (***Ivan Ivanov****, T. Moran, J. Morais, N. Desjardins, M. Ben Mekki*) . . . 4

Impact of spectrogram characteristics on image classification through transfer learning (***Bianca Granato***) . . . 5

Keynote 2 **6**

Machine Learning Analysis of Dynamical and Chaotic Processes (***Edward Ott***) . . . 6

Short talk session 2 brought to you by ORS: **7**

Dragonfly, a deep learning platform for image segmentation and more (***Nicolas Piche Ph.D.***) 7

Machine Learning for Smart Nano-electronic Biosensors (***Mohamed OUQAMRA*** *Delphine BOUILLY*) 8

Segmentation of biological tissue micro-computed tomography images with supervised deep learning (***Ksenia Kolosova****, Rui Tahara, Geoffroy Noel, Luc Mongeau, Paul Wiseman*) 8

Decoding the Drosophila gap gene network with an autoencoder (***Juliette Lavoie****, Adrien Henry and Paul Francois*) 9

Learning to work efficiently: Using neuroevolutionary strategies for reinforcement learning on classical thermodynamic systems (***Christopher Beeler****, Uladzimir Yahorau, Rory Coles , Kyle Mills, Steve Whitelam, Isaac Tamblyn*) 9

Keynote 3 **10**

Time-varying communities in brain graphs (***Pierre Bellec***) 10

Short talk session 3: **11**

Equity and Representation in Physics and STEM Fields (***Talia Martz-Oberlander***) . . . 11

Machine learning applications to gravitational physics (***Tim Whittaker****, Huan Yang, William East, Luis Lehner*) 11

Short talk session 4: **12**

The importance of realism in image-based deep neural network classifications of galaxy interactions (***Connor Bottrell****, Maan Hani, Hossen Teimoorinia, Sara Ellison, Mallory Thorp , Jorge Moreno, JPaul Torrey, Chris Hayward,*) 12

Visualising Future Extreme Climate Events at Homes with CycleGANs (***S. Karthik Mukkavilli*** *; Sasha Luccioni ; Victor Schmidt ; Yoshua Bengio*) 12

Keynote 4 **14**

Learning Protein Structure and Function (***Guillaume Lamoureux***) 14

Invited speaker **15**

Materials Informatics: the 4th paradigm (***Hong Guo***) 15

Posters **16**

Analyzing Hubble Space Telescope Data with Gaussian Processes (***Taylor Bell***, *Nikolay Nikolov, Nicolas Cowan, Joanna Barstow, Travis Barman, Ian Crossfield, Neale Gibson, Thomas Evans, David Sing, Heather Knutson, Tiffany Kataria, Joshua Lothringer, Bjrn Benneke, Joel Schwartz*) 16

Deep-Physics: Intuitive Physics of Dynamic Environments using DNNs (***Alexandre Lavoie***, *James Kary*) 16

Learning density functional theory mappings with extensive deep neural networks and deep convolutional inverse graphics networks. (***Kevin Ryczko***, *David Strubbe, Isaac Tamblyn*) 17

Attack and defence in cellular decision-making: lessons from machine learning (***Thomas J. Rademaker***, *Emmanuel Bengio, Paul Franois*) 17

Encoding and Decoding Physics Information in DNNs (***Taoli Cheng***) 18

Definign conformational states of proteins using dimensionality reduction and clustering algorithms (***Eugene Klyshko***, *Sarah Rauscher*) 18

Discovering Radio Frequency Interference using Machine Learning Methods (***Rory Coles***, *Stephen Harrison, Dave Del Rizzo, Tim Robishaw*) 18

Using Deep-learning to Characterize Gravitational Microlensing Events (***Dang, Lisa***, *Luger, Rodrigo, Cortado, Gabriella, Lam, Casey, Ishitani, Stela, Euteneuer, Esther, Penny, Matthew, Yee, Jennifer, Foreman-Mackey, Dan*) 19

Developing software to Automatically process nanometer-scale micrographs using Artificial Intelligence (***Chris Anderson***, *Jacob Klien, Laurent K. Belend, Colin Judge,Heygaan Rajakumar*) 19

Accelerometry-based inference of constrained motions (***Chris Isaac Larnder***) 20

Automated Discovery of Partial Differential Equations (*Zachary Vernec, Andy Hoang-Cao,* ***Joel Trudeau***) 20

Towards Deep Learning in Continuous-time Quantum Monte Carlo (***Simon Verret***) 21

Stochastic thermodynamics of aggregate-label learning (***Maximilian Puelma Touzel***, *Sebastian Goldt*) 21

Neural Network to Emulate Numerical Simulations of the Sun and Infer Synthetic Plasma Motions at the Photosphere (***Benoit Tremblay***, *Thierry Roudier, Jean-Franois Cossette, Michel Rieutord,Alain Vincent*) 21

A New Stochastic Local Clustering algorithm (***Hadi Papei***, *Yang Li*) 22
Neural Network Predictions of Cosmic String Locations in CMB Maps (***Razvan Ciuca*** *&*
Oscar Hernandez) 22

Workshop Outlook

Beginner's Workshop

May 6-7th
14:00-16:00
Ballroom A

Evgeny Naumov
Calcul Quebec

The Beginners Track workshop introduces foundational statistical learning ideas with examples using Python and Scikit-Learn. Participants will be introduced to the problem of binary classification and the basics of statistical learning decision surfaces, overfitting, regularization and ensemble methods using interactive online notebooks with editable code. On the second day, participants will work independently to build a model for a physics supervised regression problem.

Intermediate Workshop

May 6-7th
14:00-16:00
Salle du Parc

Jean-Philippe Reid
Element AI

The intermediate workshop WILL feature hackathon-style project-based active learning. It will introduce the datasets (specifically chosen to appeal to a physics audience) and associated methods for building interpretable and accurate classifiers. From there, participants under guidance of mentors will delve into the data and work together in building machine learning classifiers!

Keynote 1

May 6th
9:00-9:45
Ballroom A

Generative Models as AI-Assistants for NISQ Hardware

Roger Melko

Department of Physics and Astronomy, University of Waterloo

The quantum wavefunction presents the ultimate "big data" problem in physics. When a large number of qubits are entangled in a quantum computer, the resulting complexity presents a daunting challenge for any computational strategy seeking to characterize the underlying quantum state. Recently, a new computational toolbox based on modern machine learning techniques has been rapidly adopted into the field of quantum physics. In this talk, I will discuss how generative modelling with restricted Boltzmann machines can be used to combat the complexity of the quantum wavefunction. I will demonstrate successful reconstruction of the state of a cold-atom quantum computer, and discuss prospects for generative modelling in future applications in the design and characterization of Noisy Intermediate-Scale Quantum (NISQ) hardware.

Short talk session 1 brought to you by Tandem Launch:

AI + Physics : United to protect your privacy online

May 6th
10:15-10:25am
Ballroom A

Alexandre Desilets-Benoit
Tandem Launch

Advertisers are spying on us as we browse the web. They know who we are, what we do and where we are. Physics and AI meet where traditional technologies fail to protect us. This talk will cover how Contxtful leverages MEMS embedded in off-the-shelf devices to bridge the gap between us and our smart devices without breaching our privacy.

(Quantum) Machine Learning Applied to Disordered Systems

May 6th
10:25-10:35
Ballroom A

D. Bernardi, A. Cheng, D. Ittah, D. Lopez, and A. Vuong, and **M. Hilke**
Department of Physics McGill University

Machine learning has emerged as a powerful tool to classify images through training of a deep convolution neural network. An input image is then associated to an image category. Quantum machine learning extends this idea to either a quantum input, a quantum network, a quantum output, or a combination of thereof. In this work we study the case of an input, which stems from a quantum problem (the local density of states of a disordered system), a classical neural network and classical outputs. The outputs we consider are properties of the underlying random potential as well as conductance. More specifically, we considered the case of a two-dimensional electron gas in a magnetic field subjected to a random potential characterized by amplitude, impurity shape, depletion depth and roughness. We use an iterative Greens function method to evaluate the local density of states and conductance and train the neural network to recognize various properties of the system with different disorder configurations. Using multiple neural regression layers at the output allows us to predict the potential characteristics as well as the magnetic field, energy and localization properties, simply from the local density of states with over 95% confidence. Finally, the trained neural network is applied to experimental data, obtained by scanning gate microscopy on disordered graphene at low temperatures, to determine the properties of the effective underlying disorder potential.

May 6th
10:35-10:45
Ballroom A

Adaptive Quantum State Tomography with Neural Networks

Stanislav Fort[1,2*], Yihui Quek[1*], Hui Khoon Ng[3]

[1]Stanford University; [2]Google AI, [3]Yale-NUS College, Singapore, [*] Equal contributions

Quantum State Tomography is the task of determining an unknown quantum state by making measurements on identical copies of the state. Current algorithms are costly both on the experimental front – requiring vast numbers of measurements – as well as in terms of the computational time to analyze those measurements. In this paper, we address the problem of analysis speed and flexibility, introducing *Neural Adaptive Quantum State Tomography* (NA-QST), a machine learning based algorithm for quantum state tomography that adapts measurements and provides orders of magnitude faster processing while retaining state-of-the-art reconstruction accuracy. Our algorithm is inspired by particle swarm optimization and Bayesian particle-filter based adaptive methods, which we extend and enhance using neural networks. The resampling step, in which a bank of candidate solutions – particles – is refined, is in our case learned directly from data, removing the computational bottleneck of standard methods. We successfully replace the Bayesian calculation that requires computational time of O(poly(n)) with a learned heuristic whose time complexity empirically scales as O(log(n)) with the number of copies measured n, while retaining the same reconstruction accuracy. This corresponds to a factor of a million speedup for 107 copies measured. We demonstrate that our algorithm learns to work with basis, symmetric informationally complete (SIC), as well as other types of POVMs. We discuss the value of measurement adaptivity for each POVM type, demonstrating that its effect is significant only for basis POVMs. Our algorithm can be retrained within hours on a single laptop for a two-qubit situation, which suggests a feasible time-cost when extended to larger systems. It can also adapt to a subset of possible states, a choice of the type of measurement, and other experimental details.

May 6th
10:45-10:55
Ballroom A

Efficient Method for Solving the Time-Dependent Schrodinger Equation

Ivan Ivanov, T. Moran, J. Morais, N. Desjardins, M. Ben Mekki

Vanier College, Mathematics Department

It is known that time-dependent quantum systems can be formulated entirely in terms of an ensemble of quantile (Bohmian) trajectories without any reference to a wave-function. The ensemble of trajectories is computed from a non-linear partial differential equation. The numerical algorithm that we use for propagating the trajectories ensemble explicitly preserves unitarity and does not have superfluous reflections at the boundaries when the simulation region is bounded. Moreover, our method requires no fixed grids and no interpolations. Valuable insights into the dynamics of the quantum state are also gained, because the trajectories show how the dynamical process unfolds. Because the translation between the trajectory-based description and the wave-function description of a quantum system can easily be performed analytically, this method can potentially be of interest to the wider research community. We also explore the possibilities of using neural networks to perform these computations much faster and at a lower memory cost, which can prove useful for many-body problems.

Impact of spectrogram characteristics on image classification through transfer learning

May 6th
10:55-11:05am
Ballroom A

Bianca Granato

Quantitative Life Sciences, McGill University

Image classifiers are machine learning algorithms that sort images into pre-specified categories. These images could be everyday objects [1], cancer cells [2] or spectrograms [3]. Most state-of-the-art image classifiers are learning-based and require large amounts of data for training [4]. Furthermore, algorithms are often developed anew and are tailored to specific data or tasks. Through transfer learning, however,it is possible to apply a pre-trained algorithm to a related data while reducing the time and labeled data requirements[5]. For example, a classifier called AlexNet[6] was the winner of the object classification challenge of the ImageNet Large-Scale Visual Recognition Challenge in 2012[1]. AlexNet was trained to categorize objects,animals and scenery and has since been made publicly available. Even though it was not trained to classify spectrograms, previous studies have had success in transfer learning AlexNet to categorize spectrograms[7, 8] . However, the impact of image noise and spectrogram settings on classification performance has not been investigated in detail.Using previously published methods to simulate micro-Doppler signatures[9, 10] i.e. shifts in radar waves caused by moving objects, I generated spectrograms with different signal-to-noise ratios, windowing settings and colour mappings. AlexNet was first fine-tuned by training it on 15 images per micro-Doppler signature category (5 learning, 10 validation) and tested in a withheld data set. Algorithm performance was assessed through overall accuracy and with the extended Matthews Correlation Coefficient[11].

1. O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg and L. Fei-Fei, International Journal of Computer Vision 115 (3), 211-252 (2015).
2. A. A. Cruz-Roa, J. E. Arevalo Ovalle, A. Madabhushi and F. A. Gonzlez Osorio, presented at the Medical Image Computing and Computer-Assisted Intervention MICCAI 2013, Berlin, Heidelberg, 2013 (unpublished).
3. J. Dennis, H. D. Tran and H. Li, IEEE signal processing letters 18 (2), 130-133 (2011).
4. O. Boiman, E. Shechtman and M. Irani, (unpublished).
5. S. J. Pan and Q. Yang, IEEE Transactions on knowledge and data engineering 22 (10), 1345-1359 (2010).
6. A. Krizhevsky, I. Sutskever and G. E. Hinton, Advances in neural information processing systems (2012).
7. S. Zhang, S. Zhang, T. Huang and W. Gao, IEEE Transactions on Multimedia 20 (6), 1576-1590 (2018).
8. T. zseven, Applied Acoustics 142, 70-77 (2018).
9. Y. Yang, J. Lei, W. Zhang and C. Lu, (unpublished).
10. V. C. Chen, F. Li, S. S. Ho and H. Wechsler, IEEE Transactions on Aerospace and electronic systems 42 (1), 2-21 (2006).
11. J. Gorodkin, Computational Biology and Chemistry 28 (5), 367-374 (2004).

Keynote 2

May 6th
11:15-12:00
Ballroom A

Machine Learning Analysis of Dynamical and Chaotic Processes

Edward Ott

Department of Physics University of Maryland

In this talk we consider a situation where one is interested in gaining understanding the dynamics of a chaotically time evolving system through access to time series measurements that depend on the evolving state of the system. Using examples, we show that machine learning is an extremely effective tool for accomplishing this important and useful task. We also present theoretical ideas explaining the dynamical basis of how a machine learning system is able to do this.

Short talk session 2 brought to you by ORS:

Dragonfly, a deep learning platform for image segmentation and more

May 6th
13:00-13:10
Ballroom A

Nicolas Piche Ph.D.

Object Research Systems, Inc. (Montreal)

Image segmentation often requires significant time if it is done manually, especially when it comes to tomographic image stacks of 3D objects. AI-based automatic segmentation is promising but the user may be intimidated by the learning curve. Dragonfly is made for everyone to use machine learning or neural network, free of burden, for tasks from image segmentation to super resolution. Based on TensorFlow and Keras, Dragonfly is pre-builtwith popular convolutional neural networkarchitectures (e.g., UNET, FCDenseNet, PSPNet, etc.), which allow users forquickly creating their own models and training with their own image data sets. The trained models behave like regular image filters, so they can be easily previewed and applied to an entire new data set, or shared among colleagues. A basic workflow is as following: 1). create or load a neural network, 2). use Dragonflys integrated segmentation tools on select reference slices, 3). use those results as the training data to train the network, 4). use the trained network to segment the rest of the same experimental image stack and subsequent slices. Advanced users can alsotake advantage of built-in tools for drafting new networks or edit the activation functions and other node behavior in existing models. Moreover,Dragonfly users can code directly in Python with the Keras API or import pre-existing Keras models for direct integration. Dragonfly is a commercial product but it is free for non-commercial users.

May 6th
13:10-13:20
Ballroom A

Machine Learning for Smart Nano-electronic Biosensors

Mohamed OUQAMRA Delphine BOUILLY

Universite de Montreal

Nanoelectronic sensors based on field effect transistors (FETs) have demonstrated their ability to detect molecules with ultra-high sensitivity and specificity. In fact, carbon nanotube-based devices have been recently exploited to probe molecular dynamics at a single-molecule scale, an approach named smFET ("single-molecule" FET). This detection method, without any labeling of the target molecules, and strictly based on electrostatic interactions between the molecule and the transducer, allows a temporal resolution below the microsecond, offering the possibility to monitor in real time the transitions between different conformational states in biochemical reactions, such as enzymatic catalysis. A major challenge in extracting kinetics from smFET data relies both on the complexity of the molecular interactions that change between transient and steady state conformations, but also on the resulting signals, especially their non-stationary and multi-source nature. Markovian state models are widely used to model such smFET signals but require restrictive assumptions and a priori knowledge of the likely kinetics model, which are hardly met in real experiments. We propose an alternative approach based on machine learning for signal processing to retrieve the hidden dynamic network of the molecular transition states. The first learning algorithm consists in uniquely labeling each data point in the signal with a form of digital tattoo enabling to track any changes over time in the recorded traces. The second learning algorithm is a compression of the signals to minimize the potential redundancy of their information content. To model unknown kinetics, and thus to lower the false positive rate of the event detection, an unsupervised k-medoid partitioning of the compression patterns leads to idealized traces. Finally, as a last stage a decision-aid tool automatically selects the most fitting sequence of states and transitions, as the best multi-state model of the molecular dynamics.

May 6th
13:20-13:30
Ballroom A

Segmentation of biological tissue micro-computed tomography images with supervised deep learning

Ksenia Kolosova, Rui Tahara, Geoffroy Noel, Luc Mongeau, Paul Wiseman

Department of Physics & Chemistry, McGill Univeristy

Segmentation of biological images is important for creating realistic computational models from tissue images. The complexity of biological features often leads to failure of conventional first- pass segmentation methods such as simple thresholding, requiring application of more involved segmentation techniques to minimize manual segmentation time. This project explores automated segmentation using supervised deep learning with U-Net architectures as an avenue for segmenting soft tissue in high-resolution micro-computed tomography images of human and animal larynges. The project uses Dragonfly, an image processing software created by Object Research Systems, as a manual segmentation interface to create labeled training data, and as a front end for deep learning manipulations. We will present our current results and describe future directions to improve model quality, reduce manual segmentation requirements, and extract additional information from images.

Decoding the Drosophila gap gene network with an autoencoder

May 6th
13:30-13:40
Ballroom A

Juliette Lavoie, Adrien Henry and Paul Francois
Department of Physics McGill University

The gap gene network is a crucial part of the drosophila embryo development. It contains the positional information of the anterior-posterior axis that will be used to position the pair-rule genes in stripes along the embryo. Our goal is to understand how the network encode this information. We want to build an accurate model that is also interpretable. We use an autoencoder to reduce the dimensionality of the system to build our model. The autoencoder is trained on randomly shuffled concentration data of the four gap genes at different positions on the A-P axis and different times in nuclear cycle 14. By extracting the latent space of the autoencoder, we show that the dynamic of the system is encoded in two dimensions. Examining the data in this reduced space helps to gain insights on the mechanism of the network. Then, we use bayesian statistics to create maps that predict the position only from the 2d manifold. These predictions are made with a 1% error. This is the same positional error found by previous experiments on actual cells. Furthermore, from those maps, we predict the position of eve stripes in mutant embryos. The autoencoder helps us to find an interpretable representation of the network that still keeps all its important features.

Learning to work efficiently: Using neuroevolutionary strategies for reinforcement learning on classical thermodynamic systems

May 6th
13:40-13:50
Ballroom A

Christopher Beeler[1], Uladzimir Yahorau[1], Rory Coles[1] , Kyle Mills[1], Steve Whitelam[2], Isaac Tamblyn[1,3,4]
[1]University of Ontario Institute of Technology,, [2]Lawrence Berkeley National Laboratory, [3]University of Ottawa, [4]National Research Council of Canada

We show that reinforcement learning can also be used to solve classical problems in thermodynamics. Using a reinforcement learning method based on genetic algorithms, our software agent can learn to reproduce thermodynamic cycles without prior knowledge of physical laws. We have created a simulated learning environment which models a simple piston, where an agent can activate thermodynamic processes. With this method, we were able to optimize an artificial neural network based policy to maximize the thermal efficiency for several different cases. Depending on the actions available to the agent, different known cycles emerged, including the Carnot, Stirling, and Otto cycles. Importantly, we show an example of how reinforcement learning can be used to aid scientists in finding solutions to problems that have yet to be fully explored. In one of the heat engine environments, we introduced a non-adiabatic process which caused the engine to lose energy. In this case, the agent produced, what is to the best our knowledge, the best solution for the problem.

Keynote 3

May 7th
9:00-9:45
Ballroom A

Time-varying communities in brain graphs

Pierre Bellec
CRIUGM & University de Montreal

Brain activity measured with functional magnetic resonance imaging (fMRI) is dominated by neural communities with tightly synchronised fluctuations. Brain communities can easily be detected when averaging data collected across many subjects, but the common view is that tens of minutes of acquisitions are required to establish a reliable mapping for an individual subject. I will first present new work showing that, using unsupervised cluster analysis on short time windows of spontaneous activity (mns), it is possible to identify distinct communities associated with a single brain area over time. Some of these community "states" are more common than others and, despite being defined from short time windows, their spatial distribution is highly reproducible. In a second part of the talk, I will extend these results on task-based functional data from the Human Connectome Project, recorded while subjects performed one of 21 tasks designed to activate specific brain communities. I will show that it is possible to "decode" which task a subject is performing from functional brain images with excellent accuracy, using graph convolutional networks on short time windows (10 sec). Performance of brain decoding remains good even using a single volume of brain activity, lasting about 720 ms. Taken together, these results suggest that it is feasible to track precisely the dynamics of brain communities at the full temporal resolution offered by fMRI.

Short talk session 3:

Equity and Representation in Physics and STEM Fields

May 7th
9:45-10:05
Ballroom A

Talia Martz-Oberlander
Department of Physics McGill Univeristy

While some Science, Technology, Engineering, and Math (STEM) fields have seen a growth in gender diversity over the past decades, certain fields, predominantly within physics and computer sciences, lack the insights women, racialized people, or "outsiders", can contribute. This presentation from McGill Women in Physics (WiP) examines contemporary statistics and delves into some reasons why these fields are still so inequitable. By examining factors causing underrepresentation, we begin learning why certain thoughts, language, and behaviour we all take part in teach minority identities we don't belong in science and how we can instead adopt more welcoming practices. The WiP Outreach Program was developed to inspire students not societally encouraged to pursue physics or STEM to consider what careers in this area can offer them and society.

Machine learning applications to gravitational physics

May 7th
10:05-10:15
Ballroom A

Tim Whittaker, Huan Yang, William East, Luis Lehner
University of Waterloo

Over the last few years, gravitational wave detections have enabled a new window into our universe. Not only has it been a fundamental verification of the theory General Relativity but it has enabled the possibility to study objects such as neutron stars with a new perspective. Large banks of simulated waveforms are required for the detection of gravitational waves. These banks of simulated waveforms are obtained by numerically solving Einsteins field equations which is computationally costly. Meanwhile, machine learning has been gaining a large amount of attention for its abilities to classify and generate data. A natural question is then is there a way to ease the cost of generating waveforms for detection using machine learning? In this talk we will discuss the applications of generative models in gravitational physics. Specifically, we will present a conditional variational autoencoder trained to model binary neutron star gravitational waves.

Short talk session 4:

May 7th
10:45-10:55
Ballroom A

The importance of realism in image-based deep neural network classifications of galaxy interactions

Connor Bottrell[1], Maan Hani[1], Hossen Teimoorinia[2], Sara Ellison[1], Mallory Thorp [1], Jorge Moreno[3], JPaul Torrey[4], Chris Hayward[5],

[1]University of Victoria; [2]NRC [3]Pomona, [4] Florida, [5] CCA

Observations and theoretical predictions alike show that mergers transform galaxies from changes in AGN activity, star-formation rates, and gas metallicity distributions to angular momenta and morphologies. Putting these findings in an evolutionary context requires large galaxy interaction samples and the ability to connect these changes to specific stages in a merger. However, characterization of galaxy interactions through traditional methods are challenging: visual classification is subjective; quantitative morphologies are sensitive to image quality and redshift; and spectroscopic pair confirmations are notoriously incomplete. In this talk, I will discuss ways that more complete samples and more accurate merger stage classifications can be obtained by combining the hydrodynamical simulations, synthetic observations, and deep learning. Specifically, I train convolutional neural networks using synthetic observations from a suite of hydrodynamical merger simulations run with the FIRE-2 model. The simulations offer an important advantage over real observations – ground truth target classes for interaction stage. Synthetic observations of the simulations are produced with various levels of realism (e.g. projected stellar mass and dust-inclusive radiative transfer) to answer the following questions: (1) What is gained/lost by making the images more realistic? (2) How realistic do the images need to be in order to achieve high performance in classifying mergers in realistic images? (3) How sensitive is the network performance to image quality? (4) Is a network that is trained on synthetic observations for one survey easily transferable to another survey of higher or lower image quality?

May 7th
10:55-11:05
Ballroom A

Visualising Future Extreme Climate Events at Homes with CycleGANs

S. Karthik Mukkavilli ; Sasha Luccioni ; Victor Schmidt ; Yoshua Bengio

Montreal Institute for Learning Algorithms (MILA)

We present a project that aims to generate images that depict accurate, vivid, and personalized outcomes of climate change using Cycle-Consistent Generative Adversarial Networks (CycleGANs). By training our CycleGAN model on images of houses before and after extreme flood events, we learn a mapping that can then be applied to images of locations that have not yet experienced these events. This visual transformation is paired with inland and sea level flood model predictions assessing likelihood and type of climate-related events in the long term (50 years) in order to bring the future closer in the viewers mind. The eventual goal of our project is to enable individuals to make more informed choices about their climate future by creating a more visceral understanding of the effects of climate change, while maintaining scientific credibility by drawing on climate model projections. The author39;s will also strive to seek input from recent physics breakthrough39;s like Edward Ott39;s

work on simulating wildfire chaos integrating PDEs and GANs into our system framework, as it could be part of our future work.

Keynote 4

May 7th
11:15-12:00
Ballroom A

Learning Protein Structure and Function

Guillaume Lamoureux

Department of Chemistry and Center for Computational and Integrative Biology (CCIB), Rutgers University, Camden, New Jersey

Proteins are the main building blocks of living organisms. To understand any biological process in detail, one needs to know the structure, function, and dynamics of all proteins involved. Despite considerable advances in proteomics, structural and functional data are available only for a small fraction of all known proteins and an even smaller fraction of all *possible* proteins, which includes those that were not observed during evolution but that are accessible through protein engineering. To alleviate this severe data imbalance, computational methods based on sound physical principles are needed. In this talk, I will present our recent efforts at developing unified sequence-to-structure-to-function models based on deep neural networks. Following an end-to-end learning approach, these models aim to predict how proteins fold and interact with one another by discovering the data representations most useful for the tasks of predicting function from structure and structure from sequence. I will conclude by discussing some of the new research avenues opened by deep learning approaches to biomolecular data.

Invited speaker

Materials Informatics: the 4th paradigm

May 7th
13:15-14:00
Ballroom A

Hong Guo

Center for the Physics of Materials and Department of Physics, McGill University, Montreal Canada

Materials informatics (MI) may be considered the 4th paradigm of scientific inquiry, in addition to experimental, theoretical and computational approaches. MI is made possible by the universal access to abundant scientific data, assisted by advances in software and machine learning (ML) to analyze the data. For materials problems with specific designing goals, physics-based indicators (or assumptions) are necessary to help narrowing down the informatics search. In this talk I shall present materials property discovery by MI + ML, including 2D ferromagnets, solid state electrolytes, molecules for OLED, and high Tc superconductors. We conclude that backed by theory and first principles simulation, and eventually by experimental verification, MI + ML is a very efficient approach for property discovery of materials.

Posters

May 6-7th
16:00-18:00,
Ballroom B

Analyzing Hubble Space Telescope Data with Gaussian Processes

Taylor Bell, Nikolay Nikolov, Nicolas Cowan, Joanna Barstow, Travis Barman, Ian Crossfield, Neale Gibson, Thomas Evans, David Sing, Heather Knutson, Tiffany Kataria, Joshua Lothringer, Bjrn Benneke, Joel Schwartz

McGill Space Institute, Department of Physics

The characterization of extra-solar planet (exoplanet) atmospheres requires extremely precise observations. Ground-based telescopes suffer from many noise sources due to the Earths atmosphere, so many exoplanet observations use space-based telescopes such as the Hubble Space Telescope (HST). We used HST to observe the star WASP-12 at ultraviolet to optical wavelengths when its extremely hot, Jupiter-mass exoplanet passed behind the star. This allows us to search for the faint dimming of the system when we no longer see the light reflected by the planet. While HST is highly precise, there are many effects which lead to strongly correlated noise with an amplitude 10x larger than the exoplanetary signal we hope to measure. For example, HSTs orbit causes the Sun to shine on different parts of the telescope throughout the spacecrafts 95 minute orbital period. This causes thermal flexure of the optical system, leading to variations in the throughput of the telescope. The functional form of these variations is unknown and certainly complex, but the induced noise appears to vary smoothly and consistently. Traditional decorrelation techniques assume the induced fluctuations vary polynomially with auxiliary variables, but this method has been shown to bias the results and underestimate uncertainties. Instead, one can use a Gaussian process regression which assumes no functional form for the variations, instead making the minimal assumption that the function is smooth with some white noise. Combining a Gaussian process with a Markov chain Monte Carlo, we analyzed our HST observations and found that the planet reflects less than 10% of the incoming light, making it as black as fresh asphalt.

May 6-7th
16:00-18:00,
Ballroom B

Deep-Physics: Intuitive Physics of Dynamic Environments using DNNs

Alexandre Lavoie, James Kary

Vanier College

Despite revolutionary advancements in Artificial Intelligence, General AI remains a stagnant field of research. Even if Neural Networks have mastered a multitude of individual tasks, current models remain weak solutions of mimicking actual human learning. This paper attempts to break away from Narrow AI by focusing on recreating the innate human perception and understanding of the world. More specifically, this paper explores the concept of intuitive physics by proposing Deep-Physics. Using numerical data from physics models, Deep-Physics is trained to replicate humans intuitive understanding of physics. Given initial conditions of objects, Deep-Physics generates a realistic step-by-step approximation of dynamics of environments. The simplicity of the model allows Deep-Physics to be a potential alternative to real-time physics engines. Using a Deep-Physics model could be a step

into a more efficient and intuitive machine learning, eventually leading to the ultimate goal of General AI.

Learning density functional theory mappings with extensive deep neural networks and deep convolutional inverse graphics networks.

May 6-7th
16:00-18:00,
Ballroom B

Kevin Ryczko, David Strubbe, Isaac Tamblyn

University of Ottawa, Department of Physics

In this work, we show that deep neural networks (DNNs) and extensive DNNs (EDNNs) can be used in conjunction with Kohn-Sham density functional theory (KS-DFT) for two-dimensional electron gases in simple harmonic oscillator and random potentials. Using calculations from the Octopus real-space DFT code we show that EDNNs can learn the mappings between the electron density and exchange, correlation, external, kinetic and total energies simultaneously. Our results hold for local, semi-local, and hybrid exchange-correlation functionals. We then show that the external potential can also be used as input for an EDNN when predicting the aforementioned energy functionals, bypassing the KS scheme. Additionally, we show that EDNNs can be used to map the electron density calculated with a local exchange-correlation functional to energies calculated with a semi-local exchange correlation functional. Lastly, we show that deep convolutional inverse graphics networks can be used to map external potentials to their respective self-consistent electron densities. This work shows that EDNNs are generalizable and transferable given the variability of the potentials and the fact that EDNNs can scale to an arbitrary system size with an O(N) computational cost.

Attack and defence in cellular decision-making: lessons from machine learning

May 6-7th
16:00-18:00,
Ballroom B

Thomas J. Rademaker[1], Emmanuel Bengio[2], Paul Franois[1]

[1]Dept. of Physics, McGill University, [2]Dept. of Computer Science, McGill University

Machine learning algorithms are sensitive to meaningless (or "adversarial") perturbations. This is reminiscent of cellular decision-making where ligands (called "antagonists") prevent correct signalling, like in early immune recognition. We draw a formal analogy between neural networks used in machine learning and models of cellular decision-making (adaptive proofreading). We apply attacks from machine learning to simple decision-making models, and show explicitly the correspondence to antagonism by weakly bound ligands. Such antagonism is absent in more nonlinear models, which inspired us to implement a biomimetic defence in neural networks filtering out adversarial perturbations. We then apply a gradient-descent approach from machine learning to different cellular decision-making models, and we reveal the existence of two regimes characterized by the presence or absence of a critical point. The critical point causes the strongest antagonists to lie close to the threshold. This is validated in the loss landscapes of robust neural networks and cellular decision-making models, and observed experimentally for immune cells. For both regimes, we explain how associated defence mechanisms shape the geometry of the loss landscape, and why different adversarial attacks are effective in different regimes. Our work connects evolved cellular decision-making to machine learning, and motivates the design of a general theory of adversarial perturbations, both for in vivo and in silico systems.

May 6-7th
16:00-18:00,
Ballroom B

Encoding and Decoding Physics Information in DNNs

Taoli Cheng

MILA, Universit de Montreal

In the past few years, deep neural network models have been proved to be able to do a great job in High Energy Physics. A lot of interesting architectures have been brought into the market. We present a study on a physics-inspired architecture used for jet identification at the Large Hadron Collider, which makes use of the natural tree structure of jet clustering process. And based on that, we present how visualization and interpretability study can reveal encoded physics information within DNNs. Moreover, we discuss how to further utilize the encoded information in a general scope.

May 6-7th
16:00-18:00,
Ballroom B

Definign conformational states of proteins using dimensionality reduction and clustering algorithms

Eugene Klyshko[1], Sarah Rauscher[1,2]

[1]Department of Physics, University of Toronto, [2]Chemical and Physical Sciences, University of Toronto Mississauga

Molecular dynamics (MD) simulations of proteins produce large data sets - long trajectories of atomic coordinates - and provide a representation of the sampling of a given molecules structural ensemble. A deep quantitative analysis using advanced machine learning techniques is a means to interpret MD trajectories. To visualize the conformational space of the molecule and properly identify conformational states, we suggest combining clustering methods and dimensionality reduction algorithms. We investigate different choices of features to represent individual structures, clustering algorithms, similarity metric, and methods to assign the number of clusters.

May 6-7th
16:00-18:00,
Ballroom B

Discovering Radio Frequency Interference using Machine Learning Methods

Rory Coles, Stephen Harrison, Dave Del Rizzo, Tim Robishaw

University of Ontario Institute of Technology

While many steps have been taken around radio observatories to reduce Radio Frequency Interference, that which gets through often goes undetected until it shows up in the data. Monitoring for interference as it is happening requires a large number of labour-hours to be feasible and so we look to automate this process with Machine Learning methods. The goal is to obtain a real-time RFI monitoring system to assist workers on site in the mitigation of RFI at the source instead of in the data. As some RFI is persistent, we want to build a representation of the RFI scene, allowing us to identify sources that are novel. We show how Deep Neural Networks can be used to assist with RFI detection around an observatory site, and consider methods to identify patterns in the detections.

Using Deep-learning to Characterize Gravitational Microlensing Events

May 6-7th
16:00-18:00,
Ballroom B

Dang, Lisa[1], Luger, Rodrigo[2], Cortado, Gabriella[2], Lam, Casey[3], Ishitani, Stela[4], Euteneuer, Esther[5], Penny, Matthew[6], Yee, Jennifer[7], Foreman-Mackey, Dan[2]

[1]McGill University, [2]Flatiron Institute, [3]UC Berkeley,[4]Catholic University of America,[5]Heidelberg University,[6]Ohio State University,[7]Harvard Center for Astrophysics

Gravitational lensing is an astronomical effect predicted by Einsteins theory of general relativity in which the trajectory of light rays passing close to massive bodies bend due to the curvature of space-time. Today, gravitational lensing is used by astronomers as an observational technique to discover and study objects that are very far away or that emit very little light, such as dark matter, black holes, and planets. I will focus on gravitational microlensing, which is a time variable gravitational lensing of a distant source by an intervening mass where multiple images are created but not resolved. In other words, as two unrelated objects (a source of light and a dark lens) come in and out of alignment along the line of sight of an observer, the lens will bend the light from the source resulting in a magnification and demagnification of the source as seen from a distant observer.This allows us to detect lenses and modelling this effect allows us to characterize them, however, producing such models to fit our observation can be computationally expensive. Modelling the effect of a microlensing event can be broken down in two steps: 1) modelling the lensing properties of the lens and 2) modelling the relative trajectory between the source and the lens in the course of your observations. In January, I attended a Microlensing Hack session, during which participants were encouraged to tackle unresolved problems in the field or provide improvements to existing limitation. As an attempt speed up the modelling process, we used deep-learning to train a neural network using precomputed magnification maps of lenses and to predict magnification maps given some lens properties. During this talk, I will share the results of our hack project.

Developing software to Automatically process nanometer-scale micrographs using Artificial Intelligence

May 6-7th
16:00-18:00,
Ballroom B

Chris Anderson[1], Jacob Klien[1], Laurent K. Belend[1], Colin Judge[2],Heygaan Rajakumar[2]

[1]Queens University [2]Chalk River Laboratories

The present project aims at processing Transmission Electron Micrographs (TEM) of neutron-irradiated Ni-Superalloys using an Artificial Neural Net (ANN) which is a recently developed method of machine learning approach, to automatize the identification of irradiation-induced nano-structures. This method works by learning features of the micro-structural defects during a training phase with previously identified data and then applying those learned features to previously unseen images. The project aims to reduce the need for manual processing of images and to allow for statistically significant quantities of data to be processed. Focus is on the training methods used to identify the bubbles as well as the challenges that exist in this unique environment. The model has shown promising results with classifications of 80% of total bubbles captured and 85% of total bubble surface area. The current state of the model is discussed along with future steps and present outstanding challenges.

May 6-7th
16:00-18:00,
Ballroom B

Accelerometry-based inference of constrained motions

Chris Isaac Larnder

Physics Department, John Abbott College

The reconstruction of motion from accelerometer data is, in the general case, difficult. One problem is the noise-amplification drift inherent to double-integration techniques. Another fundamental problem is the seperation of the proper acceleration into gravitational and coordinate-acceleration components. Einsteins equivalence principle precludes any hardware-level solutions to this latter problem, recently-announced atom interferometry technologies notwithstanding.

For motions subject to mechanical constraints, however, appropriate exploitation of the reduction in the degrees of freedom could lead to significant improvements in or complete elimination of these problems. In such an approach, a set of differential equations provides a parametrised characterisation of both the constraints imposed on the physical object and the geometrical relationship between the object and the sensor. By evaluating the convergence of multi-dimensional minimization algorithms, we expect to automatically infer the type of constraint at play. A corollary of the successful detection of such constraint signatures in the accelerometer data is the recovery of reliable estimates for all parameters in the model, including, for a broad class of constraints, state variables for the reduced-coordinate degrees of freedom.

A methodical exploration of some of these opportunities is the subject of a recently-awarded 3-year FRQNT grant whose scope and goals will be outlined in this presentation. Although our study begins with models having a small number of parameters, the development of such techniques promises broad applicability as a low-level classification layer in machine-learning architectures involving more complex physical systems. Recognizing that traditional ML systems suffer from inflexibility and poor transfer to related domains, there is increasing interest in hybrid systems that incorporate domain-specific knowledge. Different deep networks specialized for distinct application domains could eventually re-use the same physics layer whose output is also human-interpretable, reducing development costs and improving system comprehensibility.

May 6-7th
16:00-18:00,
Ballroom B

Automated Discovery of Partial Differential Equations

Zachary Vernec, Andy Hoang-Cao, **Joel Trudeau**

Dawson College, Physics

This project is applies methods described in recent work from Samuel H. Rudy, Nathan Kutz, et al. on the identification of parametric partial differential equations from data using grouped sparse regression methods. We test these methods on data generated from a larger set of equations than the few canonical examples given in the paper, that is, equations with different number of terms and orders, different coefficients, and different parameters. We measure the performance of the methods along those examples. The data is generated directly using finite-difference methods over a sufficiently fine grid.

Towards Deep Learning in Continuous-time Quantum Monte Carlo

May 6-7th 16:00-18:00, Ballroom B

Simon Verret
MILA

Strongly correlated electrons are an important frontier of condensed matter physics. Research in this field often uncovers new properties of matter, with possible applications in energy production, new electronic and thermic devices, new sensors, etc. This poster reproduces results from a 2017 paper by Nagai et al., entitled Self-learning Monte-Carlo method: continuous-time algorithm. As in the paper, we use machine learning to improve Monte Carlo sampling for impurity solvers used to simulate materials with strongly interacting electrons. The long-term goal is to explore more sophisticated generative models, such as variational autoencoder (VAE) and generative adversarial networks (GAN)

Stochastic thermodynamics of aggregate-label learning

May 6-7th 16:00-18:00, Ballroom B

Maximilian Puelma Touzel[1], Sebastian Goldt[2]
[1] MILA, Universite de Montreal,[2]IPhT, CEA Saclay

How do learning agents improve their temporally-extended actions over repeated trials of a task when the feedback signal they receive is delayed and intermittent? This temporal credit assignment problem is at the core of many sequence learning tasks in machine learning and behavioural neuroscience. The multi-spike tempotron is a recently proposed classifier that solves the problem in a variety of tasks by using only the recent aggregate feature count. It outperforms in sample complexity conventional reinforcement learning by orders of magnitude. Here, we consider a student-teacher framework, wherein the classifier is tasked with matching the spikerate of another already trained multi-spike tempotron. Focusing on biologically plausible learning rules that exploit the correlation between pre and post synaptic activity, we show that the student neuron also achieves high performance in this task. Formulating the dynamics of alignment of the students weights with those of the teachers allows for quantifying the stochastic thermodynamic efficiency of this learning rule for classification of spiking circuit activity by single neurons.

Neural Network to Emulate Numerical Simulations of the Sun and Infer Synthetic Plasma Motions at the Photosphere

May 6-7th 16:00-18:00, Ballroom B

Benoit Tremblay[1], Thierry Roudier[2], Jean-Franois Cossette[1], Michel Rieutord[2],Alain Vincent[1]
[1]Universit de Montral, Canada [2]IRAP, Universit de Toulouse III, Toulouse, France

Eruptive events of the Sun, which often occur in the context of flares, convert large amounts of magnetic energy into emission and particle acceleration that can have significant impacts on Earth's environment. Satellites and ground-based observatories probe the Sun's photosphere and atmosphere and are key in studying solar activity. The wealth of data available has been instrumental in investigating physical features relevant to the onset of flares through statistical analyses and machine-learning algorithms. Meanwhile, numerical models have attempted to bridge the gap between the physics of the solar interior and such observations. However, there are relevant physical quantities that can be modelled but that cannot be directly measured and must be inferred. For example, direct measurements of plasma motions at the photosphere are limited to the line-of-sight component. Recently, neural network computing has been used in conjunction with numerical models of the Sun to be able to recover the full velocity vector in photospheric plasma of the Quiet Sun (i.e. in the absence of significant magnetic activity). We used satellite observations as input in a fully convolutional neural network to generate instantaneous synthetic plasma motions, i.e. plasma motions that reflect the physics of a model but

are made to look as if they were observed by a specific instrument. A parallel technique could then be invoked to eventually be able to derive the plasma velocity vector maps of solar active regions (i.e. regions of significant magnetic activity) and, by extension, other physical quantities of interest that can not yet be measured directly at the photosphere or anywhere else in the solar atmosphere.

May 6-7th
16:00-18:00,
Ballroom B

A New Stochastic Local Clustering algorithm

Hadi Papei[1], Yang Li[2]

[1]Department of Physics and Astronomy, Western University, London, Ontario, N6A 3K7, Canada
[2]Department of Mathematics and Statistics, University of Minnesota Duluth, Duluth, MN 55812, USA

We propose a novel stochastic local clustering algorithm to find clusters around any given member of a network represented as a graph. Determining the structure within datasets has become an important aspect of research in both science and industry. Most clustering algorithms require the entire dataset to probe its structure; however, some datasets, e.g., WWW, are too large to be stored entirely or even be fully accessible. Our local algorithm is appropriate for such data when it is only practical to access a subset, or if we are mainly interested in finding the structure around a data point of interest. Unlike deterministic algorithms that find a subset of data as a cluster, our stochastic algorithm calculates the probability that any other members of the network could be within the cluster of the initial seed member. This feature makes the algorithm suitable for the types of data in which a sharp boundary between clusters is not expected, e.g., social media networks or gene expression patterns. Other advantages of this stochastic algorithm over deterministic algorithms are the ability to both evaluate membership uncertainties and explore hierarchical structures within the data. The performance of the algorithm is illustrated for synthetic and real data.

May 6-7th
16:00-18:00
Ballroom B

Neural Network Predictions of Cosmic String Locations in CMB Maps

Razvan Ciuca & Oscar Hernandez

Department of Physics McGill University

We use convolutional neural networks to look for signatures of cosmic strings in numerical simulations of comic microwave background (CMB) temperature maps. This application of deep learning is unique in that we search for a very subdominant signal. String-induced temperature fluctuations realistically account for much less than 1% of the total CMB power, and in this regime standard deep learning wisdom fails. Still, we successfully trained networks that were able to accurately predict string locations and the value of the string tension. In this talk, I will discuss the importance of cosmic strings and their signal in the CMB, explain our neural network architecture, and present our results. I will then discuss the pitfalls we encountered and why more classical methods failed. I will end by describing some useful tricks for training deep networks on similar datasets.
