

Opportunity Title: Mathematical Intuition in Neural Discovery and Understanding (MINDU)

Opportunity Reference Code: ICPD-2019-18

Organization Office of the Director of National Intelligence (ODNI)

Reference Code ICPD-2019-18

How to Apply **Create and release your Profile on Zintellect** – Postdoctoral applicants must create an account and complete a profile in the on-line application system. **Please note: your resume/CV may not exceed 2 pages.**

Complete your application – Enter the rest of the information required for the IC Postdoc Program Research Opportunity. The application itself contains detailed instructions for each one of these components: availability, citizenship, transcripts, dissertation abstract, publication and presentation plan, and information about your Research Advisor co-applicant.

Additional information about the IC Postdoctoral Research Fellowship Program is available on the program website located at: <https://orise.orau.gov/icpostdoc/index.html>.

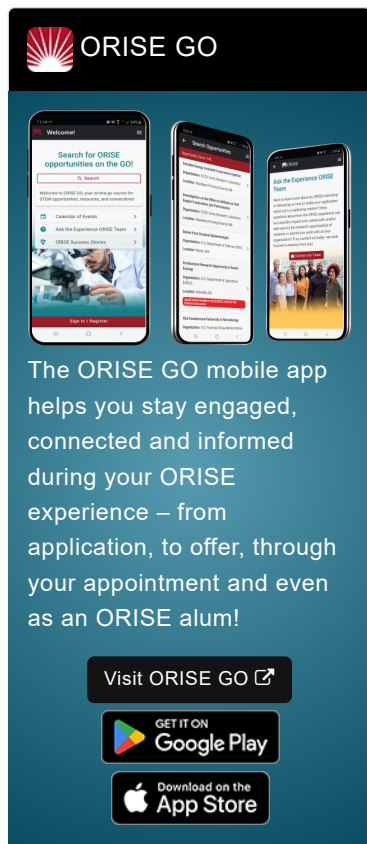
If you have questions, send an email to ICPostdoc@orau.org. Please include the reference code for this opportunity in your email.

Application Deadline 3/4/2019 5:00:00 PM Eastern Time Zone

Description **Research Topic Description, including Problem Statement:**

- In recent years, there has been a tremendous growth in research in machine learning and artificial intelligence. This growth is spurred by trends toward large curated data sets, availability of General Purpose Graphical Processing Unit (GPGPU)-enabled cloud environments and free collaborative ecosystems for rapidly distributing and testing new software. While the growth in Machine Learning and Artificial Intelligence (ML/AI) continues to expand, there has not been matching progress in understanding the implications of new ML/AI applications. In particular, recent deep learning techniques have shown tremendous power in addressing computer vision and human language technology problems, but are shown to succumb to broadly engineered adversarial attacks.
- Much of the research on explainable ML/AI has focused on the security of ML/AI techniques, and on better understanding and characterizing failure modes. The work has generally looked at the problem from an experimental, or statistical learning theory point of view. There has been little research that taps into previous deep mathematical theory that is relevant in this domain.
- In particular, recurrent neural networks (RNNs) such as Long, Short-Term Memory (LSTM) models may be viewed as dynamical systems preserving an invariant measure, or stationary processes. Note that well-studied Markov Processes (or Hidden Markov Models, HMMs) are stationary. There is a long history of research on measurable dynamics, including that of von Neumann, Kolmogorov, Shannon, Ornstein, Furstenberg, and more recently Fields Medalists: Lindenstrauss, Avila, Mirzakhani and Venkatesh. In 2015, Fields Medalist Terence Tao posted a 77-page preprint on the Navier-Stokes Equations; while falling short of solving the Millenium Problem, did shed new light on the possibility of a counter-example. On page 15, Tao writes:

"It is worth pointing out, however, that even if this program is successful, it would only demonstrate blowup for a very specific type of initial data (and tiny perturbations thereof), and is not necessarily in contradiction with the belief that one has global regularity for most choices of initial data (for some carefully chosen definition of "most", e.g. with overwhelming (but not almost sure) probability with respect to various



ORISE GO

The ORISE GO mobile app helps you stay engaged, connected and informed during your ORISE experience – from application, to offer, through your appointment and even as an ORISE alum!

Visit ORISE GO

GET IT ON Google Play

Download on the App Store

Opportunity Title: Mathematical Intuition in Neural Discovery and Understanding
(MINDU)

Opportunity Reference Code: ICPD-2019-18

probability distributions of initial data). However, we do not have any new ideas to contribute on how to address this latter question, other than to state the obvious fact that deterministic methods alone are unlikely to be sufficient to resolve the problem, and that stochastic methods (e.g. those based on invariant measures) are probably needed."

- There are many questions that may be explored with respect to Deep Learning (DL) models from a mathematical perspective. DL models demonstrate a sensitivity to initial conditions, both in the training regimes, and during run-time classification. Local minima distribute throughout the parameter space. It would be interesting to understand how local minima spread throughout the parameter space? Which local minima track closely the output by an absolute minimum, and how do local minima distribute as models increase in size, and architectures change? How do regularizations effect convergence to local minima or potentially to an absolute minimum? Is there any advantage in introducing regularization (that is convexity constraints)? Potentially viewing DL models in a larger hyperspace will yield insights. Could there be ever evolving or constantly learning networks which stay ahead of potential adversarial attacks? Could a better mathematical understanding of Deep Neural Networks (DNNs) lead to better network design or more efficient models as well as improved convergence and generalizations of models?

Example Approaches:

- One approach is to apply the theory of function approximation or related principles from harmonic analysis. One example can be seen by the research of Mhaskar and Poggio, *Deep vs. Shallow Networks: an Approximation Theory Perspective*.
- Other recent research has pointed to the strengths and limitations of an important framework, Generative Adversarial Networks (GANs). GANs can be used to boost training and generalization of DNN, as well as create realistic synthetic data.
- When the data to be modeled has a temporal aspect, other areas come into play such as non-parametric statistics, time series modeling, statistical physics, topological and measurable dynamics, statistical learning theory, time series prediction and econometrics. There has been little research applying principles from pure and applied mathematics to better understand deep neural networks.
- There could be advantages to exploring a universal representation of stationary processes and how these representations might be used for prediction or learning features for distinguishing various temporal phenomenon. Examples include RNNs or possibly other mathematical models such as adic transformations defined on Bratelli-Vershik diagrams (that is a type of directed acyclic graph).
- Another approach might explore the renormalization group. For example, see the 2014 preprint *An Exact Mapping between the Variational Renormalization Group and Deep Learning* by P. Mehta and D.J. Schwab.

Key Words:

Machine Learning; Artificial Intelligence; Deep Neural Networks; Deep Learning; Explainable AI; Generative Adversarial Networks; Adversarial Attacks; Non-parametric Bayesian Networks; Probabilistic Programming; Neural Chunking; Probabilistic Graphical Models; Spectral Theory; Harmonic Analysis; Operator Theory; Function Approximation; Dynamical Systems; Lie Algebras;

Opportunity Title: Mathematical Intuition in Neural Discovery and Understanding (MINDU)

Opportunity Reference Code: ICPD-2019-18

Bratteli-Vershik Diagrams; Control Theory; N-gram modeling; Hidden Markov Models; Model Capacity; Statistical Learning Theory; PAC Learning; Vapnik-Chervonenkis Dimension; Stochastic/Stationary Processes; Entropy; Fisher Information; Information Bottleneck; Statistical Mechanics; Renormalization Group.

Qualifications **Postdoc Eligibility**

- U.S. citizens only
- Ph.D. in a relevant field must be completed before beginning the appointment and within five years of the application deadline
- Proposal must be associated with an accredited U.S. university, college, or U.S. government laboratory
- Eligible candidates may only receive one award from the IC Postdoctoral Research Fellowship Program.

Research Advisor Eligibility

- Must be an employee of an accredited U.S. university, college or U.S. government laboratory
- Are not required to be U.S. citizens

Eligibility Requirements

- **Citizenship:** U.S. Citizen Only
- **Degree:** Doctoral Degree.
- **Discipline(s):**
 - **Chemistry and Materials Sciences** ([12](#) 👁)
 - **Communications and Graphics Design** ([6](#) 👁)
 - **Computer, Information, and Data Sciences** ([16](#) 👁)
 - **Earth and Geosciences** ([21](#) 👁)
 - **Engineering** ([27](#) 👁)
 - **Environmental and Marine Sciences** ([14](#) 👁)
 - **Life Health and Medical Sciences** ([45](#) 👁)
 - **Mathematics and Statistics** ([10](#) 👁)
 - **Other Non-Science & Engineering** ([5](#) 👁)
 - **Physics** ([16](#) 👁)
 - **Science & Engineering-related** ([1](#) 👁)
 - **Social and Behavioral Sciences** ([28](#) 👁)