

Center for Science of Information

NSF Science and Technology Center

Prestige Lecture 2021



Alexander Gray, Ph.D.
Vice President
Foundations of AI, IBM

Monday, January 25
2:00pm EST

ONLINE

<https://www.soihub.org>
See under Events & Activities



Center for Science of Information

Logical Neural Networks: Toward Unifying Statistical and Symbolic AI

Abstract:

Recently there has been renewed interest in the long-standing goal of somehow unifying the capabilities of both statistical AI (learning and prediction) and symbolic AI (knowledge representation and reasoning). We introduce Logical Neural Networks, a new neuro-symbolic framework which identifies and leverages a 1-to-1 correspondence between an artificial neuron and a logic gate in a weighted form of real-valued logic. With a few key modifications of the standard modern neural network, we construct a model which performs the equivalent of logical inference rules such as modus ponens within the message-passing paradigm of neural networks, and utilizes a new form of loss, contradiction loss, which maximizes logical consistency in the face of imperfect and inconsistent knowledge. The result differs significantly from other neuro-symbolic ideas in that 1) the model is fully disentangled and understandable since every neuron has a meaning, 2) the model can perform both classical logical deduction and its real-valued generalization (which allows for the representation and propagation of uncertainty) exactly, as special cases, as opposed to approximately as in nearly all other approaches, and 3) the model is compositional and modular, allowing for fully reusable knowledge across tasks. The framework has already enabled state-of-the-art results in several problems, including question answering.

Bio:

Alexander Gray serves as VP of Foundations of AI at IBM, and currently leads a global research program in Neuro-Symbolic AI at IBM. He received AB degrees in Applied Mathematics and Computer Science from UC Berkeley and a PhD in Computer Science from Carnegie Mellon University. Before IBM he worked at NASA, served as a tenured Associate Professor at the Georgia Institute of Technology, and co-founded and sold an AI startup in Silicon Valley. His work on machine learning, statistics, and algorithms for massive datasets, predating the movement of “big data” in industry, has been honored with a number of research honors including multiple best paper awards, the NSF CAREER Award, selection as a National Academy of Sciences Kavli Scholar, and service as a member of the 2010 National Academy of Sciences Committee on the Analysis of Massive Data. His current interests generally revolve around the injection of non-mainstream ideas into ML/AI to attempt to break through long-standing bottlenecks of the field.

