NJIT

Departments of Computer Science & Data Science

Fundamental Problems in AI: Transferability, Compressibility and Generalization

Tomer Galanti
Center for Brains, Minds, and Machines at MIT

Hosted by Hai Phan + Yiannis Koutis

Abstract:
In this talk, we delve into several fundamental questions in deep learning. We start by addressing the question, "What are good representations of data?" Recent studies have shown that the representations learned by a single classifier over multiple classes can be easily adapted to new classes with very few samples. We offer a compelling explanation for this behavior by drawing a relationship between transferability and an emergent property known as neural collapse. Additionally, we explore why certain architectures, such as convolutional networks, outperform fully-connected networks, providing theoretical support for how their inherent sparsity aids learning with fewer samples. Lastly, I present recent findings on how training hyperparameters implicitly control the ranks of weight matrices, consequently affecting the model's compressibility and the dimensionality of the learned features.

Additionally, I will describe how this research integrates into a broader research program where I aim to develop realistic models of contemporary learning settings to guide practices in deep learning and artificial intelligence. Utilizing both theory and experiments, I study fundamental questions in the field of deep learning, including why certain architectural choices improve performance or convergence rates, when transfer learning and self-supervised learning work, and what kinds of data representations are learned with Stochastic Gradient Descent.

Bio:
Tomer Galanti is a Postdoctoral Associate at the Center for Brains, Minds, and Machines at MIT, where he focuses on the theoretical and algorithmic aspects of deep learning. He received his Ph.D. in Computer Science from Tel Aviv University and served as a Research Scientist Intern at Google DeepMind's Foundations team during his doctoral studies. He has published numerous papers in top-tier conferences and journals, including NeurIPS, ICML, ICLR, and JMLR. His work, titled "On the Modularity of Hypernetworks," was awarded an oral presentation at NeurIPS 2020.