Robin.ly Dec 11, 2019 5 min read

[NeurIPS 2019 Highlight] Martin Schrimpf @ MIT: Brain-Like Object Recognition with Recurrent ANNs

Updated: Jan 15

This episode is an interview with Martin Schrimpf Ph.D. student in Brain and Cognitive Sciences at MIT BCS. He shared highlights from his paper Brain-Like Object Recognition with High-Performing Shallow Recurrent ANNs, which was accepted for oral presentation at NeurIPS 2019 conference.

This paper proposed a quantitative collaboration between neuroscience and machine learning by representing a brain score that allows you to compare models with the brain and developing a brain-like model, CORnet-S, which transforms deep networks into much more shallow networks with recurrence. This research will help build architectures and networks that are more like the brain and improve energy efficiency for computing.

Martin's main interest is in bridging Machine Learning and Neuroscience with a focus on building deep neural network models of the brain's ventral stream that are more human-like in their behavior as well as their internals. His previous work includes research in computer vision at Harvard, and natural language processing and reinforcement learning at Salesforce.

Interview with Robin.ly:

Robin.ly is a content platform dedicated to helping engineers and researchers develop leadership, entrepreneurship, and AI insights to scale their impacts in the new tech era.

Subscribe to our newsletter to stay updated for more CVPR interviews and inspiring talks:

Paper At A Glance

Deep convolutional artificial neural networks (ANNs) are the leading class of candidate models of the mechanisms of visual processing in the primate ventral stream. While initially inspired by brain anatomy, over the past years, these ANNs have evolved from a simple eight-layer architecture in AlexNet to extremely deep and branching architectures, demonstrating increasingly better object categorization performance, yet bringing into question how brain-like they still are. In particular, typical deep models from the machine learning community are often hard to map onto the brain's anatomy due to their vast number of layers and missing biologically-important connections, such as recurrence. Here we demonstrate that better anatomical alignment to the brain and high performance on machine learning as well as neuroscience measures do not have to be in contradiction. We developed CORnet-S, a shallow ANN with four anatomically mapped areas and recurrent connectivity, guided by Brain-Score, a new large-scale composite of neural and behavioral benchmarks for quantifying the functional fidelity of models of the primate ventral visual stream. Despite being significantly shallower than most models, CORnet-S is the top model on Brain-Score and outperforms similarly compact models on ImageNet. Moreover, our extensive analyses of CORnet-S circuitry variants reveal that recurrence is the main predictive factor of both Brain-Score and ImageNet top-1 performance. Finally, we report that the temporal evolution of the CORnet-S "IT" neural population resembles the actual monkey IT population dynamics. Taken together, these results establish CORnet-S, a compact, recurrent ANN, as the current best model of the primate ventral visual stream. [presentation slides]