Research Statement

Seeking to understand how to cultivate learning is my passion as both an educator and as a researcher. Although humans and machines still operate in quite different ways, it is my belief that the underlying principles that facilitate effective learning are fundamentally "platform independent". Learning is one of the capabilities that makes humans profoundly effective, and giving analogous capabilities to machines will enable them to become much more helpful in assisting and empowering humans.

Machine learning is already a significant component in such applications as automated medical diagnosis, facial recognition, fraud detection, product recommendation, automated trading, personal digital assistants, and self-driving vehicles. Much of my research revolves around a brain-inspired artificial learning model called a deep neural network (DNN). While I have been studying them, DNNs have surged in popularity, establishing "deep learning" as a highly recognized buzz-word in scientific communities. This broad interest has triggered significant investment from companies such as Google, Facebook, IBM, and Amazon, which has led to rapid advancement and expansion of the frontier in this domain. DNNs have achieved super-human accuracy in many image recognition competitions and machine learning benchmarks.

However, until machines can learn as effectively as humans, my research goals are not yet complete. My research focuses on improving the learning capabilities of machines in four key areas of high potential for major advancements: collaborative learning, extracting intrinsic representations from high-dimensional data, learning from temporal data, and combining various learning methods for intelligent agents and robots.

Collaborative Learning: Humans not only learn from their individual experiences, but also through interacting with each other. My research has achieved improvements in learning through augmenting model diversity in ensembles [C2] and through combining deep learning methods with collaborative filters [C6, J4]. One method I invented aligns neural network topologies such that their weights can be fused with minimal loss of accuracy [C8]. Recent results from my lab that are still pending publication demonstrate that we can apply model fusion to learn in a distributed manner. Experiments show that our models can learn indirectly, and demonstrate capabilities for which they were never explicitly trained. This might lead, for example, to effective methods for learning to cure diseases by mining hospital data without requiring the sensitive patient data to ever be aggregated.

In order to strengthen the impact of my work in machine learning, I developed an open source toolkit of machine learning algorithms called *Waffles* [J2]. This toolkit provides an excellent mechanism for me to distribute my contributions where many researchers will see and use them. It also gives significant positive exposure to my laboratory and the University of Arkansas. It contains more than 100,000 lines (about 2,000 printed pages) of source code implementing hundreds of different algorithms for a wide variety of machine learning applications. It includes algorithms and tools for classification, regression, collaborative filtering, non-linear dimensionality reduction, clustering, Bayesian graphical models, ensemble methods, data visualization, sparse matrix processing, optimization, reinforcement learning, and deep neural network learning. Two companies have invited me to speak to their engineers about my toolkit, and I have received more than 20 unsolicited letters of thanks from researchers around the world who use my code in their projects.

Extracting Intrinsic Representations from High-dimensional Data: One of the greatest cognitive strengths that humans demonstrate is the ability to understand the things they observe. For example, when looking at a picture, humans do not merely understand it as a 2D array of pixel values, but can describe the objects they see in the picture. A similar task is performed by non-linear dimensionality reduction algorithms, which produce brief summaries, or low-dimensional encodings, of otherwise complex data.

My research thrust is to enable machines to learn in ways that are not currently possible. Thus, my work with methods that extract meaningful representations from high-dimensional data [C1, C3, C4, C6, C11, C12, J1, J3, J4] have not merely focused on improving round-trip accuracy, but have largely focused on learning meaningful representations from unconventional data, or in unconventional situations. In particular I have focused on applying these methods to problems that change over time, such as robotics challenges, because I view these to be more analogous with the challenges that humans have already learned to handle effectively. It is my belief that the key to enabling machines to operate effectively under diverse conditions is to first enable them to understand the observations that they make.

Learning from Temporal Data: Much of the existing machine learning research focuses on addressing static problems, such as recognizing stationary faces or written characters. Although much success has been found in these areas, we live in a world that is constantly changing. Since time tends to have a confounding effect on learning, several of my research efforts have focused on mitigating this challenge [C4, C5, C9, C13, J2, J5, J6]. In particular, my works lean towards consuming data as high-dimensional observations that are expected to change over time. This non-traditional approach has potential to open the way for models suitable for use in complex changing environments, such as those associated with real-world factories, robots, and autonomous vehicles.

Intelligent Agents and Robots: One of the grand objectives associated with machine learning is to combine the various learning methods into an effective cognitive architecture. This is an endeavor in which I am very interested [C7], in part because it serves as a sort of idea generator for my on-going research. When I come across a student in need of a new research idea, I consider where machines currently fall short in learning capabilities in comparison with humans, and I inevitably find a challenge in need of further study.

My work in addressing novel learning challenges has led to several advances that improve the effectiveness of intelligent agents and robots [C3, C4, C10, J3, J7]. I foresee both significant growth and funding opportunities in this area, as more activities that previously relied on human labor begin to transition toward automation. Some indications that this is already beginning to happen include the emergence of self-driving cars, automated drone-based delivery systems, increased use of robotics in surgery, automated manufacturing, rover-based exploration, recent successes in automated laundry folding, and the increased prevalence of robots in the home for vacuuming, mopping, and gardening.