Mike Gashler's Research Statement

I am deeply interested in the problem of giving machines a human-like ability to understand, to think and to reason. I am persuaded that the ability to learn is one of the primary factors that enables humans to operate effectively with such a diversity of problems, and that advancing our understanding of computational learning techniques will ultimately lead to giving machines more of the abilities that are currently unique to humans. Machine learning has also found significant utility in more immediate applications, such as data mining, recommendation systems, robotics, etc. These immediate applications have been a boon for research in machine learning because they have provided both necessary funding and motivation to thoroughly explore the field. I believe that the best way to achieve my ultimate goal of giving machines human-like capabilities is to find stepping-stone advances that can be firmly established with immediate applications.

My specific research interest is focused on automatic methods for extracting meaning from images. Unlike typical computer vision approaches, which tend to use human experts to design domain-specific metrics that convert images to meaningful features, my work focuses on using machine learning approaches that automatically learn to recognize what is important in images. Although learning-based approaches are not yet as effective as expert-designed approaches, learning-based methods offer certain obvious advantages. Perhaps more significantly, I believe that learning-based methods are poised to soon overtake expert-designed methods in effectiveness, which will result in a significant shift in the approaches used to analyze images. This shift toward learning methods is likely to lead to many new avenues for making significant advances in the capabilities of machines for understanding images. My immediate goal is to help to push existing learning approaches past this tipping point.

While I have focused primarily on images, the methods I work with are also applicable to other high-dimensional data, including documents, audio, large sensor arrays, etc. My recent research has found several techniques that improve the effectiveness of manifold learning algorithms, and also make them applicable to a wider scope of problems. For example, some of my recent work has discovered the utility of using estimated state in the process of training recurrent models of dynamical systems. This provides both a new approach for performing black-box system identification, and a quantitative approach for evaluating the quality of state estimates.

My research agenda for the near future begins with an exploration of using non-linear PCA, a manifold learning technique based on artificial neural networks, with the application of black-box system identification. I expect to find that this approach is even better-suited for this task than methods with which I have found success in the past. Further, this will help to prepare for my subsequent research plans. Looking somewhat farther out, I strongly suspect that deep network approaches will lead to the next several leaps in advancing this field. Already, advances in parallel hardware, such as CUDA, are making it possible to train much deeper networks than could previously be trained within a reasonable amount of time. It appears that it will soon become reasonable to train very deep networks using a life-long-learning approach in an on-going simulated environment. I am hopeful that such a system may automatically develop the underlying structures necessary to facilitate adept capabilities in recognition and vision with real-world problems.

High-dimensional data from visual, tactile, and other types of sensor arrays often contains a wealth of intrinsic meaning and state information that humans are very capable of comprehending. Ultimately, I intend for the systems that I help to build to be able to accurately extract this intrinsic information, such that machines may be able to operate with a level of understanding and comprehension comparable to that with which humans currently operate.