Adaptive Regression and Personalized Routing via Interactive Learning

Nick Terry
Ph. D. Student
Industrial & Systems Engineering
University of Washington

Abstract: Machine learning algorithms which adjust their own learning strategy over time can outperform static approaches in many scenarios. We discuss two such paradigms: adaptive methods which can reactively accommodate violations of modeling assumptions, and active methods which proactively guide the sequential collection of training data. As an example of adaptive learning, we present a computationally efficient approach to regression based on recursively splitting Gaussian processes, along with applications in robotics and power generation. In contrast, we give an overview of our recent work in active learning. This work is motivated by the problem of personalized routing, where a traveler's individual preferences determine the cost of traversing each edge in a graph. We describe an efficient approximation algorithm for learning the traveler's preferences in large search spaces and computing a route between a fixed origin and destination which is near-optimal with respect to the traveler's implicit objective function.

Bio: Nick Terry is a Ph.D. student in the Industrial and Systems Engineering department at the University of Washington, working with Dr. Youngjun Choe (ISE) and Dr. Kevin Jamieson (CSE). His research focuses on statistical machine learning and sequential experiments with applications in engineering and operations management. He is currently serving as an Air Force Research Laboratory Scholar to study active machine learning. Prior to his Ph.D. studies, Nick spent two years working as a supply chain engineering consultant for Fortna, Inc. He holds a Bachelor of Science degree from the Georgia Institute of Technology and a Master of Science degree from the University of Washington, both in Industrial Engineering.