Hristo Paskov Teaching Statement

My passion for teaching comes from my passion for learning: those realization moments when the material comes together and forever changes what you know and how you think. My goal as a teacher is to provide students with the tools and courage to solve hard problems and to show them the broader picture, connections, and beauty in mathematics and science. Like my research, my teaching approach relies on leveraging structure in the material. This includes distilling concepts down to their essence and presenting them from the viewpoint that best appeals to each student's background. Having taken classes my entire life I have a deep appreciation for which class format pair best with each subject, and I have been lucky to have taken some exceptionally taught classes that I will use as models for my own.

Experience

I draw on teaching experience from working as a teaching assistant, mentoring high school students on science projects, and the exemplary classes that I have taken. I served as a teaching assistant in 2014 and 2015 for Stanford's Mining Massive Datasets (CS246), and I used this opportunity to hone my teaching approach. My responsibilities included leading weekly recitation and grading homework. I also chose to teach lectures on linear algebra and optimization and designed numerous homework and final questions. During recitation I focused on teaching students how to explore mathematical structure. Since the majority were computer scientists, I appealed to their familiarity with modularity and system design in showing them how to break hard problems into simpler, tractable pieces. This approach helped numerous students who were struggling with the material and I developed a consistent and large following in both years.

I also mentored my brother, Ivan, from 2012-2014 on a high school science project that applied machine learning to personalized medicine. Teaching a high school student enough about machine learning and rigorous research to beat the state of the art in cancer drug efficacy prediction presented unique challenges that allowed me to refine my teaching approach. The first year was spent learning about basic machine learning and experimental methodologies in a project-based format that replicated existing work to provide hands on experience. I carefully balanced the material between narrowly focusing on specific methods (so as not to be overwhelming) and providing abstractions and context. By the second year Ivan was ready to perform research in earnest and he spent the summer implementing and applying new machine learning algorithms under my guidance. The project resulted in a scientific publication and it has won numerous awards including fifth place in the Siemens Competition and finalist in the Intel Talent Search. I am now actively mentoring my youngest brother on a science project applying machine learning to DNA analysis.

Finally, I have taken several classes which were particularly effective at delivering material. One such course is Stanford's Convex Optimization (EE364A) taught by Professor Stephen Boyd. It covers a large number of topics in convex optimization with surprising depth by carefully using abstraction to highlight the structure of convex problems and then narrowing in on concrete examples to further explore concepts and provide familiarity. The professor's presentation style is encouraging and it inspires students to boldly tackle hard problems; I believe that this kind of encouragement is particularly helpful in mathematics. The class final is a 24-hour take home exam that provides students with the time to think deeply about questions that test realistic uses of convexity.

Community

I founded and organized a joint computer science and statistics reading group at Stanford for two years along with fellow advisee Will Fithian. The group's primary focus was to foster interaction between members of the two departments that were interested in machine learning. It was well attended and featured presentations from a number of prominent professors in statistics and machine learning. The group has since been subsumed by a larger artificial intelligence reading and debate group led by professors. I look forward to fostering interaction among my future communities and to inspiring new, cross-departmental ideas.

Example Courses

Undergraduate

1. *Applied Machine Learning*: A project-based course that introduces students to the basic techniques in machine learning and data analysis. The focus will be on providing students with the experience to apply these techniques and to interpret their results and behavior.
2. *Great Ideas in Mathematical Modeling*: A survey course that exposes students to a variety of ideas in statistics, linear programming, and non-linear optimization. The goal will be to provide students with an appreciation of the high-level concepts that have been revolutionary (such as Khachiyan's algorithm for linear programming) or proven essential over the ages (such as regression).
3. *Analyzing Sequential Data*: An advanced undergraduate course that covers algorithms for categorical and continuous sequential data, including applications of suffix trees and hidden Markov models. The course will emphasize algorithmic reasoning and problem solving using these techniques.
4. *The Geometry Behind Computational Complexity*: A survey course that explores different geometrical objects, their connection to optimization and algorithms, and their implications for complexity. Special attention will be paid to convexity and duality, especially spheres and cubes and their implications for separability and strong convexity, as well as polyhedra and NP-Completeness.

Graduate

1. *Machine Learning and Big Data*: Foundational course in machine learning that focuses on techniques in supervised and unsupervised learning, learning theory, regularization, and essential algorithms. The class will be steered towards large-scale data analysis and will feature a final project.
2. *Convex Optimization*: Emphasizes modeling with convexity as well as classical optimization methods for differentiable functions and their convergence analysis. Homework and final problems will challenge students and will provide additional depth and examples to the material.
3. *Large-Scale Convex Optimization*: A second course in optimization that focuses on the newest techniques for large scale problems. The course will emphasize leveraging problem structure and understanding the connections between methods using core concepts from convex analysis.