John Blitzer

Teaching Statement

> *Learning when coupled with application, is it not delightful*?
> 学而时习之, 不亦说乎?
> Confucius. The Analects. 1:1

As a computer science educator, my primary objective is to enable students to apply what they have learned. To that end, I design courses that aid students in achieving their own goals. When I teach undergraduates, I keep in mind that most do not have research and teaching as their primary professional objectives, but they will be the next generation of engineers. For these students, I design lectures and course work around specific systems that arise in industry. In contrast, most graduate students do have research as their primary professional goal, and in the tutorials and short courses I have taught for natural language processing (NLP) research audiences, my focus has been on machine learning methods that can specifically improve NLP research.

As a graduate student at the University of Pennsylvania, I served as the teaching assistant for undergraduate Numerical Linear Algebra (`http://www.math.upenn.edu/~robim/`) the first time it was taught. The topics of the course necessarily involve a significant amount of theory, but professor C.J. Taylor and I designed the lessons and projects around the engineering problems that involve direct application of that theory. For example, when the syllabus covered singular and eigenvalue decompositions, we taught lectures and designed projects around information retrieval and collaborative filtering. These are core tasks for companies like Amazon and Google, where many Penn graduates end up as engineers.

After graduating, I continued to teach as a visiting researcher at Microsoft Research Asia. At the 2008 meeting of the Association for Computational Linguistics, I taught a tutorial on semi-supervised learning (`http://ssl-acl08.wikidot.com/start`) to NLP researchers. These researchers did not come to see proofs of improved rates of convergence. Instead, I focused on semi-supervised learning techniques that can be applied to the high-dimensional, sparse problems

that we encounter in NLP. I believe this focus was one of the main reasons that out of the six tutorials taught that year, mine was the best-attended.

I expanded the course for a two day summer school on natural language processing, jointly hosted by Microsoft Research Asia and the Harbin Institute of Technology (`http://john.blitzer.com/harbin/index.html`). In addition to semi-supervised learning, my expanded course also served as an introduction to machine learning for NLP graduate students. One of my primary topics was feature design, because although it is almost never covered in machine learning courses, feature design is essential to the effective use of machine learning in NLP. The summer school was attended by 105 students from across China, and I was pleased to receive student comments such as

> *... a clear lecture about important concepts in classification and semi-supervised learning. His sense of humor and excellent Chinese also helped me a lot to understand his lecture.*
> *... good at guiding students to finding answers themselves.*
> *... one of the most enlightening tutorials I've heard on machine learning.*

Finally, one of the most important ways Ph.D. students apply what they've learned is through collaborative research with their advisors. As advisors, we must be aware of this in everything from the projects we begin with students to the papers we suggest them to read. I have already taken some roles as an advisor. At Penn, I started one student, Ted Sandler, on a project that combined his interests in spectral graph methods with an application in high-dimensional, sparse feature spaces. This led to a top conference publication. At Microsoft Research Asia, I had the good fortune to work closely with Wei Gao, a Ph.D. student in information retrieval from the Chinese University of Hong Kong. Together, we applied bilingual learning methods to web search, a topic he has continued to pursue and on which he is now writing his Ph.D. thesis. At Berkeley, I am advising two more great students on projects that I hope will lead to new research threads for them and me, and I can't wait to begin my own advisor-student relationships as a professor.

I look forward to applying the same principle of *learning coupled with application* when designing semester-long course curricula. Undergraduate courses in numerical methods, optimization, and machine learning are becoming increasingly important for software engineers. Even when such courses exist, though, they often lack the connection to the real world tasks that engineers need to solve on a daily basis. I would like to teach an undergraduate course that expands the ideas I explored as a teaching assistant for Numerical Linear Algebra – one that emphasizes industrial applications in areas like data mining and natural language processing. In a similar vein, there are

many important concepts in online and semi-supervised learning that are applicable to everyday research in data-hungry fields like NLP, and a seminar or graduate class that teaches them needs to focus on the benefits to natural language processing research.

Helping students achieve their goals by applying what they learn is the driving force behind what I did as a teaching assistant and tutorial instructor. I'm excited to take that philosophy forward as a professor as well, from curriculum design to Ph.D. student advising.