

At my first meeting with the Knowledge Discovery in Databases (KDD) Research Group, Dr. William H. Hsu, the KDD Director, asked me to prepare an hour-long presentation on a topic of my choosing. As a sophomore this was rather overwhelming, but I obliged and began the search for a topic. It was then 2004, and at that time, social network sites such as *Facebook*, *LiveJournal* and *MySpace* were rapidly gaining popularity. I, along with many of my peers, had joined one of them. Motivated by the rising popularity of social networks I reviewed literature on the topic and spent several weeks preparing a presentation to give to the KDD group.

I became specifically interested in the work of Dr. Lise Getoor who argued that many real-world datasets contain interrelated entities and capture relationships among these entities. I observed that link mining could be applied to social networks to predict the existence of relationships where none are annotated. The feeling that I was solely responsible for an original insight helped clarify my desire to conduct research in computer science and made me realize that this was indeed what I wanted to do with my life. Dr. Hsu was also very interested in my findings and said that he believed they could be the basis for more research. I later devoted my Honors Research project to study link mining and its broader implications.

Shortly thereafter I had the opportunity to co-author a research paper in AAI's Spring Symposium in Stanford, California. Later I co-authored and presented more mature results at the full paper session of the International Conference on Weblogs and Social Media in Boulder, Colorado. In these and other papers, we showed that, by learning a classifier based on several inter-member relationships and other features, we could predict with very high accuracy, precision and recall the existence of annotated links between members. This project included work on web mining, graph theory and machine learning, and I became fascinated by the intersection of these three fields. Not only was it a great experience to meet many people whose research I had reviewed, but it was also great to discover that many researchers were excited about our research and were planning to use our approaches in their own projects.

I realized at the conference that research is more than just an individual pursuit by a single person or group; instead, people around the world are actively pursuing complementary goals. By discovering that I was a part of this global network and that I had improved the quality and expediency of other projects by performing this research, I became even more determined to be as active in this community as possible. To that end, I plan to receive a Ph.D. and work as a tenure-track professor at a major research university. Ph.D. candidacy and a career as a professor would allow me to contribute broader impacts to computer science through service, teaching and research. I am fortunate to have already been given several such opportunities.

Service: Throughout my undergraduate and into the start of my graduate career I have been very active in several outreach initiatives and other university-wide projects. As an undergraduate I helped start the Kansas State Chapter of the Association of Computing Machinery (ACM) which is still growing steadily. I also restarted the Computer Science department's participation in the university's open house, an invaluable recruitment tool, which the now self-sustaining ACM continues to lead admirably. As a five-term student senator I have led other senators in allocating millions of dollars to fund the operation of the K-State Student Union, Lafene Health Center, Athletics Department and dozens of other vital organizations. As a senior I was appointed by the student-body president to the faculty senate, and with the help of several faculty members I was able to change the university schedule for class times to be more responsive to students' needs. Perhaps more importantly, I built strong relationships with caucus leaders in the faculty senate.

Although my term is over I am still frequently asked to attend certain hearings and work as a liaison for the student-body president to the faculty senate. Currently, I serve as the Student Director of Technology and advise the university administration on matters pertaining to university technology. During this time the university has deployed new financial and student information systems and is currently pursuing a new university-wide email and calendar system. I believe that using the Internet as a platform for the dissemination of knowledge is a noble task. Therefore, I am spearheading a movement to make certain instructional and lecture material freely available via the Internet as podcasts through the popular iTunesU service.

Teaching: As a teaching assistant for the introduction to computer science course I am able to instruct a lab with students from many disciplines and levels of expertise. I have been fortunate to receive very high marks on my teaching evaluations from my students: 4.9/5.0 in overall effectiveness. I thoroughly enjoy teaching; so much so that I have repeatedly turned down offers for a full time research assistant position because such an arrangement would prevent me from being able to teach. As it stands, I am happily able to do both.

Research: As a senior member of the KDD group I now lead a team of graduate and undergraduate researchers in the development and implementation of several projects. On my own initiative I have successfully started new research projects, such as the content extraction subgroup which employs several undergraduate assistants, and seen them through to completion. These projects have each culminated in one or more peer-reviewed publications. I am often chosen to lead in collaborative research efforts such as the *Webster* project at the Discrete Sciences Summer Institute at the University of Illinois Urbana-Champaign (UIUC) and the Web-Mining project with the Integrated Information Analysis Center at Kansas State University.

The research program I outline in my Proposed Plan of Research has the potential to introduce into the computer science literature new algorithms and methodologies unparalleled in their accuracy and speed. These advances will be adaptable and implemented to solve problems ubiquitous in computing, and their popularization could lead to progress in many other scientific and industrial fields. The algorithms and methodologies I develop are currently, and will continue to be, freely available to interested researchers; for instance, a Google-search for my name returns third-party programs that contain source code written by me. A NSF GRF will further enable me to pursue my long-term research goals by eliminating the need to spend time and effort applying and reapplying for funding each fiscal year; it will also allow me to apply the time otherwise spent scrounging for funding towards exhibitions and the dissemination of my research.

Computer science research of this type is inherently multi-disciplinary. Already I have begun collaborations with fellow text extraction researchers from Germany's Mainz University, social networking researchers from Italy's Milan Polytechnic, and machine learning experts from UIUC to name a few. Furthermore, I helped one particular researcher (and reference-letter writer), Dr. Jan Crow, design and complete experiments for her dissertation in behavioral psychology. The NSF GRF will give me the financial freedom to continue similar activities that enrich the lives of not only computer scientists, but the community at large.

I thank you each for the time you have spent evaluating this application, and with your help I look forward to continuing my work in the future.