

My undergraduate computing research spans technical work in algorithms, graph theory, multi-camera tracking, and distributed computing as well as a qualitative investigation into CS education. I played significant leadership roles in all four projects and have been a primary author of two publication efforts. My technical experience provides excellent preparation for my proposed graduate research in automated analysis of multi-camera surveillance networks.

Novice Programmer Debugging Strategies (2007-8, *Olin College*). As a sophomore, I studied the behavior of novice programmers under Dr. Matthew Jadud. I uncovered common syntactical and semantic problems by tracing the compile-by-compile progress of novices. I authored a paper that analyzed how novices used strategies such as “divide and conquer” or “move on” to debug string formatting errors and recommended intervention strategies for CS educators. Learning effective debugging strategies is not easy for novices. As a future professor I hope to apply lessons from this research to make computing a more accessible and attractive field.

Characterizing Nash Equilibria in Betweenness Centrality Network Formation Games (*Summer 2008, RIPS REU at UCLA, Sponsored by Microsoft Research Asia*). As a rising junior eager to sink my teeth into technical research, I took a position at RIPS, an REU hosted by UCLA’s Institute for Pure and Applied Mathematics. Mentored by Prof. Gunes Ercal and sponsored by Dr. Wei Chen of Microsoft Research Asia’s theory group, my team sought to apply game theory to understand equilibria that arise in ad-hoc network formation processes. We studied the Bounded Budget Betweenness Centrality (B3C) network formation game, where each network node seeks to maximize its share of network traffic (betweenness) while constrained to a uniformly fixed number of connections. Maximizing betweenness is a natural objective in many real-world network scenarios, such as a server that profits by increasing the traffic it delivers.

Our primary objective was to find and characterize Nash equilibria solutions to the B3C game. We applied analytical and simulation techniques to identify stable graphs and analyze associated structural properties such as connectivity and symmetry. Our notable breakthroughs include proving that all B3C networks at Nash equilibria must be connected and describing a novel algorithm to approximate a player’s NP-Hard best response calculation in polynomial time. We also studied approximate Nash equilibria and found that some structural families of graphs have constant asymptotic bounds on approximate equilibria despite increasing graph size.

At RIPS, I served as Project Manager, a leadership role that included the crucial task of communicating our work to a broader academic audience. During the project, I set our research agenda and reported results to our sponsor in weekly conference calls to Beijing. Additionally, I was responsible for maintaining internal momentum and ensuring that all team members found ways to contribute. After the project officially ended, I devoted my free time to submitting our work to an academic conference. I sought review from external sources and oversaw revision of our final paper. Leading this project from exploratory stages to a competitive (though unsuccessful) conference submission proved tremendously rewarding.

RIPS’ focus on international participation forced me to refine my teamwork skills to overcome geographical and cultural gaps. With two Chinese students on my team, I learned firsthand how to navigate cultural differences in group discussion and paper writing habits. In the end, I established lasting friendships with my teammates Xiao and Xiaohui as we bonded over pick-up basketball and spicy noodle dinners. I hope to pursue cross-cultural collaboration in future research to seek innovative perspectives and a broader audience for results.

Automated Topology Estimation of a Multi-Camera Surveillance Network (*Summer 2009, MIT Lincoln Laboratory*). After completing my theoretically-oriented RIPS research, I sought to round-out my undergraduate research exposure in project with an application-driven focus. As a rising senior I happily became a summer intern at MIT Lincoln Laboratory. I joined a team of three full-time staff lead by Dr. Jason Thornton tasked with assessing and improving current techniques for automated video surveillance of wide-area critical infrastructure sites.

My group's overall goal that summer was to invent an automated multi-camera tracking system. Working alone with occasional guidance, I tackled the problem of predicting which camera a moving target will appear in next after exiting a given view. This is especially difficult when views do not overlap, which often happens in real-world implementations. Seeking camera-to-camera transition prediction, I built software that could estimate the relative locations of cameras along common traffic routes. This system considers two cameras to lie on a common traffic route if a statistical pattern exists in the timing of respective entry and exit signals. To detect this pattern in camera footage, I extended Makris et al.'s cross-correlation approach (see project proposal) for inferring camera-to-camera transition time distributions so that it required only foreground motion detection data rather than a single-camera object tracker.

A major focus of my project was evaluating my system's performance on sparse and dense traffic scenarios for networks of four to five cameras. I concluded that given only minutes of footage my topology inference method could successfully infer roughly half of ground truth topological links in sparse-traffic facilities and detect some wide-gap links even in dense traffic. Most false negatives were attributable to poor background subtraction or heavy occlusion. Pushed by my group leader to prove exactly how well these estimated links improved the track handoff problem, I showed that for sparse traffic scenes my procedure could predict a given target's transition to a neighboring, non-overlapping camera with near 100% accuracy. I also showed that the predicted time window captured very few confuser objects (one on average) in addition to the target, making the correspondence decision a tractable problem.

As its primary impact, my work at Lincoln Laboratory provided government decision-makers with an encouraging assessment of the capabilities and shortcomings of current surveillance technology. More broadly, my work contributes toward efforts to produce a functional tracking system, which could dramatically improve protection of secure facilities by allowing officers to track suspects throughout the facility in real time, even across large gaps in coverage. Effective camera-to-camera route prediction will be a crucial step towards this goal.

A New Architecture for Distributed Control Systems (*Fall 2009-present, Olin College, Sponsored by IBM Research*). I serve as Project Manager for a five-member team on this final-year capstone project experience guided by Prof. Mark Sheldon and IBM Fellow Alex Morrow. Our goal is to prototype a control systems architecture that provides secure communication over public networks, assures resilient recovery on failure, and guarantees sensor data authenticity. Due to a non-disclosure agreement and active patent process, I cannot reveal further details.

Publications

Hughes, M., Jadud, M., & Rodrigo, M. "Novice Programmer Strategies for String Formatting." Under Review for the Journal of Computer Science Education.

Bei, X., Chen, W., Ercal-Ozkaya, G., Fu, X., **Hughes, M.,** & McBride, S. "On Pure and Approximate Nash Equilibria in Betweenness Centrality Games." Unpublished manuscript.