Bruce MacLeod develops cell phone technology for health care delivery in Ghana

Bruce MacLeod (Ph.D.'89)

High quality and timely health information is not readily available to pregnant women and new mothers in the Upper East region of Northern Ghana. Internet access is virtually non-existent, literacy rates are low, and travel distances to clinics can be significant. Yet the need for health information is especially important. Malaria remains the major cause of death among under-5 children, and proper use of intermittent preventative treatment during pregnancy can reduce anemia, fetal loss, and dangerous complications for the unborn child. Local beliefs and superstitions often guide women when information about best practices is not available. In the Upper East, elder female relatives encourage younger pregnant women to labor at home for as long as they can stand before going to a health facility, as a mark of strength and willpower. When women do come to a health clinic, they are given a booklet of basic health information for the pregnancy; unfortunately, most women are unable to access the information because they are illiterate.

The combination of poverty, remote villages, low literacy, and widespread disease creates a health crisis for women and children all over Africa. Lack of information, and bottlenecks in the gathering of data and transfer of information, contribute to the negative health outcomes. Although the Ghana Health System coordinates local volunteer health workers, outreach nurses, clinics, and district health centers to provide care and track the success of interventions, information gaps and challenges pervade the health system.

"Yet cell phone coverage for this remote and rural society is expanding rapidly. In Ghana, the coverage even includes GPRS capabilities which allow communication between cell phones and internet servers. Most of the women in the Upper East of Ghana have mobile phones and our baseline research found that only 16% of women had no access," says Bruce MacLeod, Software Architect on the Mobile Technology for Community Health initiative (MoTeCH), and Associate Professor of Computer Science at the University of Southern Maine.

The team's MoTeCH project attempts a systemic change in Ghana's existing health information system through integrating mobile phones into the rural health system. Pregnant women and new mothers with access to a phone can get relevant, personalized health information from MoTeCH. Nurses are given a low-cost mobile phone for entering data and receive MoTeCH messages about patients overdue for health care. The data entered by nurses is also used to automate the generation of reports for district and regional health administrators.

To enable this functionality, electronic medical records for each women and child in the study area are kept. Basic information about health service delivery and health status are maintained in these records. From this data, MoTeCH calculates the upcoming schedule of recommended care for each patient. For example, schedules of immunization shots are developed for each newborn child and upcoming clinic visit dates are developed for each pregnant woman. When care is due, MoTeCH notifies patients and health workers by mobile phone. As new information comes into the medical record, the schedule of upcoming care for a patient may be updated. While using electronic medical records in this rural setting may seem quite a far reach, it actually solves significant and long standing communication and information transfer problems in a relatively simple and low cost manner.

A pregnant woman who will be receiving MoTeCH messages

The underlying software architecture of MoTeCH evolved from a strategic decision to build on "best-of-class" software applications that have been extensively field tested in Africa, adds MacLeod. The mobile phone data collection system is based on openXData (**www.openxdata.org**). Health information about women and children is maintained in the OpenMRS medical record software application (**www.openmrs.org**). Finally, the task of delivery voice messages to patients was provided by the IntelliIVR (**www.yo.co.ug**) application. Significant project risk was eliminated because these systems have gone through multiple iterations of deployment and revisions in environments similar to those in NorthEast Ghana.

"Developing interoperability logic between these systems was one of the primary software development challenges we faced," says MacLeod. "Data collected by nurses using mobile phones needed to be uploaded into the appropriate places in a medical record. The timing of phone calls to patients using the IVR system is based on data from the electronic medical system. When patients do not attend the clinic within the recommended time frame, which can be calculated from the data in the medical record, text messages are sent to the nurse's mobile phone."

MacLeod's software development work was part of a larger collaboration between Columbia University, the Grameen Foundation, and the Ghana Health Service. The Gates Foundation funded the project with the goal of understanding how mobile phones could be used to achieve measurable health impacts. The system has been functioning since the beginning of the year in the Upper East region of Ghana. Preliminary reactions from patients receiving MoTeCH messages have been generally positive and

A community health officer entering data

nurses are becoming more proficient in using the technology. Ultimately, understanding the full impact and potential of a system like MoTeCH requires a multi-dimensional evaluation, examining the intervention effects on health workers and patients. This evaluation is underway.

"I have had the privilege of working on software systems and health problems in Africa and Asia since I was a graduate student at UMass Amherst," says MacLeod. "My first project sent me to the Sudan in 1984 to work with the Ford Foundation. That was my introduction to the simple joy of working with committed research scientists who are striving to improve health care in these rural communities."

MacLeod has over two decades of experience in developing appropriate software technology for health research in the developing world and in making complex software systems accessible to health researchers who are non-computer scientists. The MoTeCH system is the software standard for INDEPTH network of health research centers (**www.indepth-network.org**). Advised by Prof. Robbie Moll while in the department, MacLeod received an M.S. and Ph.D. in Computer Science from UMass Amherst in 1985 and 1989, respectively.

A Kassena-Nankana extended family compound

In December, The Federal Communications Commission (FCC) appointed **Henning Schulzrinne** (Ph.D. '92) as Chief Technology Officer. He will advise on matters across the agency to ensure that FCC policies are driving technological innovation, including serving as a resource to FCC Commissioners. Schulzrinne is Julian Clarence Levi Professor of Mathematical Methods and Computer Science and Professor of Engineering at the Fu Foundation School of Engineering at Columbia University. He has been an Engineering Fellow at the FCC since 2010. This fall, Schulzrinne also received the IEEE William Terry Lifetime Distinguished Service Award.

A number of CS alums and current faculty had their publications included in the recent list of SIGSOFT Top 10 Most Cited Articles. **Alexander Wolf** (Ph.D. '85) and co-author were listed #1 for "Foundations for the study of software architecture," **James Corbett** (Ph.D. '92), **Matthew Dwyer** (Ph.D. '95), and co-authors were second for "Bandera: extracting finite-state models from Java source code," **Peri Tarr** (Ph.D. '96), former Sr. Research Fellow **Stanley Sutton**, and co-authors were third for "N degrees of separation: multi-dimensional separation of concerns," Prof. **Leon Osterweil** was seventh for "Software processes are software too," and **Matthew B. Dwyer**, Professor **George S. Avrunin**, and **James C. Corbett** were eighth for "Patterns in property specifications for finite-state verification."

CIIR Alum **Giridhar Kumaran** (Ph.D. '08) and his wife Ramya Vaidyanathan announced the birth of their son Pranav Giridhar, born on September 1st.

This fall, **Mario Barrenechea** (B.S. '11) was on the winning team of the Winward Code War Contest held at the University of Colorado Boulder where he is currently a graduate student.

Mimno named 2011 CIFellow

David Mimno (Ph.D. Feb. '12) was named one of the twenty Ph.D. grads to the Computing Community Consortium (CCC) 2011 class of Computing Innovation Fellows (**cifellows.org**). For his fellowship, Mimno will hold a post-doc position at Princeton University working with mentor David Blei.

According to the CCC, "these 20 talented researchers have been competitively awarded postdoctoral positions of up to two years at academic institutions and industrial research laboratories throughout the U.S. Made possible by a $6.5 million National Science Foundation (NSF) grant to CRA - the third grant in as many years - the 2011 CIFellowships are a continuation of the highly successful effort begun in 2009 to forestall a permanent loss of research talent likely to occur as a consequence of the financial crisis and subsequent economic downturn."

Recent Computer Science Ph.D. graduates (AY 2010-2011)

Christopher Amato; *Increasing Scalability in Algorithms for Centralized and Decentralized Partially Observable Markov Decision Processes: Efficient Decision-Making and Coordination in Uncertain Environments*; (Shlomo Zilberstein, Advisor); Sept. 2010; Analytics, Modeling and Simulation Scientist, Aptima, Inc.

As agents are built for ever more complex environments, methods that consider the uncertainty in the system have strong advantages. This uncertainty is common in domains such as robot navigation, medical diagnosis and treatment, inventory management, sensor networks and e-commerce. When a single decision maker is present, the partially observable Markov decision process (POMDP) model is a popular and powerful choice. When choices are made in a decentralized manner by a set of decision makers, the problem can be modeled as a decentralized partially observable Markov decision process (DEC-POMDP). While POMDPs and DEC-POMDPs offer rich frameworks for sequential decision making under uncertainty, the computational complexity of each model presents an important research challenge. As a way to address this high complexity, this thesis develops several solution methods based on utilizing domain structure, memory-bounded representations and sampling. These approaches address some of the major bottlenecks for decision-making in real-world uncertain systems. The methods include a more efficient optimal algorithm for DEC-POMDPs as well as scalable approximate algorithms for POMDPs and DEC-POMDPs. Key contributions include optimizing compact representations as well as automatic structure extraction and exploitation. These approaches increase the scalability of algorithms, while also increasing their solution quality.

Bo An; *Automated Negotiation for Complex Multi-Agent Resource Allocation;* (Victor Lesser, Advisor); Feb. 2011; Postdoctoral Associate, Dept. of Computer Science, University of Southern California.

Automated negotiation (bargaining) is the most widely used approach for multi-agent resource allocation and it has recently received increasing attention. However, information uncertainty, existence of multiple contracting partners and competitors, agents' incentive to maximize individual utilities, and market dynamics make it difficult to calculate agents' rational equilibrium negotiation strategies and to develop successful negotiation agents which behave well in practice. This thesis is concerned with analyzing agents' rational behavior and developing negotiation strategies for a range of complex negotiation contexts. First, we consider the problem of finding agents' rational strategies in bargaining with incomplete information. We provide an algorithm based on the combination of game theoretic analysis and search techniques, which finds agents' equilibrium in pure strategies when they exist. Next, we extend the alternating-offers protocol to handle concurrent negotiations in which each agent has multiple trading opportunities and faces market competition. We provide an algorithm based on backward induction to compute the subgame perfect equilibrium of concurrent negotiation. Third, we present the design and implementation of agents that concurrently negotiate with other entities for acquiring multiple resources. Finally, we consider the problem of allocating networked resources in dynamic environment. We propose a distributed negotiation mechanism where agents negotiate over both a contract price and a decommitment penalty.

Aruna Balasubramanian; *Architecting Protocols to Enable Mobile Applications in Diverse Wireless Networks;* (Arun Venkataramani and Brian Levine, Advisors); Feb. 2011; CI Fellow, Dept. of Computer Science, University of Washington.

The goal of my thesis is to architect robust protocols that overcome disruptions and enable applications in diverse mobile networks. Mobile networks are prone to frequent disruptions and therefore cannot support most Internet applications. In this thesis, I focus on four networks that span the diverse connectivity spectrum to answer the question: What protocol design best overcomes disruptions and enables applications in a given network? I design four utility-driven protocols that tolerate disruptions in an environment by leveraging opportunism. Specifically, I present: 1) RAPID, which uses opportunistic replication to enable bulk transfer in mostly disconnected networks; 2) Thedu, which uses opportunistic prefetching to enable web search in intermittently connected networks; 3) ViFi, which uses opportunistic forwarding to enable Voice over IP (VoIP) in mostly connected mesh networks; and 4) Wiffler, which uses opportunistic augmentation to improve application performance in mostly connected cellular networks. Finally, I present a detailed evaluation of the protocols using implementation and deployment experiments on two large scale vehicular testbeds. Deployment on a real testbed shows that the protocols are practical and can be implemented in realistic usage environments. The evaluations show that the protocols significantly improve performance of applications compared to the state-of-the-art, in their respective environments.

George Bissias; *Bounds on Service Quality for Networks Subject to Augmentation and Attack;* (Brian Levine, Advisor); Sept. 2010; Algorithms Developer, Fluent Mobile.

Assessing a network's vulnerability to attack and random failure is a difficult and important problem that changes with network application and representation. We furnish algorithms that bound the robustness of a network under attack. We utilize both static graph-based and dynamic trace-driven representations to construct solutions appropriate for different scenarios. For static graphs we first introduce a spectral technique for developing a lower bound on the number of connected pairs of vertices in a graph after edge removal, which we apply to random graphs and the power grid of the Philippines. To address the problem of resource availability in networks we develop a second technique for bounding the number of nominally designated client vertices that can be disconnected from all server vertices after either edge or vertex removal (or both). Dynamic networks are modeled as disruption tolerant networks (DTNs). DTNs are composed of mobile nodes that are intermittently connected via short-range wireless radios. In the context of both human and vehicular mobility networks we study both the effect of targeted node removal and the effect of augmentation with stationary relays.

TJ Brunette; *Adaptive Balancing of Exploitation with Exploration to Improve Protein Structure Prediction;* (Oliver Brock, Advisor); May 2011; Senior Fellow, Dept. of Biochemistry, University of Washington.

One of the most significant impediments for protein structure prediction is the inadequacy of conformation space search. Conformation space is too large and the energy landscape too rugged for existing search methods to consistently find near-optimal minima. Conformation space search methods thus have to focus exploration on a small fraction of the search space. The ability to choose appropriate regions, i.e. regions that are highly likely to contain the native state, critically impacts the effectiveness of search. To make the choice of where to explore requires information, with higher-quality information resulting in better choices. Most current search methods are designed to work in as many domains as possible, which leads to less accurate information because of the need for generality. However, most domains provide unique, and accurate information. To best utilize domain-specific information, search needs to be customized for each domain. The first contribution of this thesis customizes search for protein structure prediction, resulting in significantly more accurate protein structure predictions. My results indicate that integrating the information between homologs and fragments significantly improves protein structure prediction accuracy, resulting in several proteins predicted with 1 angstrom RMSD resolution.

Bin Chen; *Improving Processes Using Static Analysis Techniques;* (Lori Clarke and George Avrunin, Advisors); Feb. 2011; Quantitative Developer, Vegasoul Capital.

Real-world processes often undergo improvements to meet certain goals, such as coping with changed requirements, eliminating defects, improving the quality of the products, and reducing costs. Identifying and evaluating the defects or errors in the process, identifying the causes of such defects, and validating proposed improvements all require careful analysis of the process. Human-intensive processes are of particular concern because they can be extremely complex and may be used in critical, including life-critical, situations. To date, the analysis support for such processes is very limited. There has been considerable success lately in using static analysis techniques to analyze hardware systems, software systems, and manufacturing processes. This thesis explores how such analysis techniques can be automated and employed to effectively analyze life-critical, human-intensive processes. We investigated two static analysis techniques: Finite-State Verification (FSV) and Fault Tree Analysis (FTA). We proposed a process analysis framework that is capable of performing both FSV and FTA on rigorously defined processes. We evaluated this framework based on the Little-JIL process definition language and employed it to analyze two real-world, human-intensive processes – an In-Patient Blood Transfusion Process and a Chemo Therapy Process. The results show that the framework can be used effectively to detect defects in such real-world, human-intensive processes.

Michael Hay; *Enabling Accurate Analysis of Private Network Data;* (Gerome Miklau and David Jensen, Advisors); Sept. 2010; Postdoctoral Associate, Dept. of Computer Science, Cornell University.

This dissertation addresses the challenge of enabling accurate analysis of network data while ensuring the protection of network participants' privacy. Massive amounts of data are being collected (facebook activity, email correspondence, cell phone records), and there is huge interest in analyzing the data, but the data is not being shared due to concerns about privacy. Despite much research in privacy-preserving data analysis, existing technologies are inadequate because they were designed for tables, not networks, and cannot be easily adapted to handle the complexities of network data. We develop several technologies to meet this challenge. First, we develop a framework for assessing the risk of publishing a network that has been "anonymized." Using this framework, we show that a small amount of background knowledge about local network structure suffices to re-identify an "anonymous" individual. This motivates our second contribution: an algorithm that transforms network structure to provably lower re-identification risk. In comparison with other algorithms, our approach more accurately preserves important features of the network topology. Finally, we consider an alternative paradigm, in which the analyst accesses data through a regulated query interface. We show that the degree sequence of a network can be accurately estimated while ensuring strong privacy protection.

Vidit Jain; *Using Context to Enhance the Understanding of Face Images;* (Erik Learned-Miller, Advisor); Sept. 2010; Scientist, Yahoo! Labs Bangalore.

Faces are special objects of interest. Developing automated systems for detecting and recognizing faces is useful in a variety of application domains including providing aid to visually-impaired people and managing large-scale collections of images. Humans have a remarkable ability to detect and identify faces in an image, but related automated systems perform poorly in real-world scenarios, particularly on faces that are difficult to detect and recognize. Why are humans so good? There is general agreement in the cognitive science community that the human brain uses the context of the scene shown in an image to solve the difficult cases of detection and recognition. We focus on emulating this approach by using different kinds of contextual information for improving the performance of various approaches for face detection and face recognition: an algorithm that employs the easy-to-detect faces in an image to find the difficult-to-detect faces in the same image and a joint probabilistic model for image-caption pairs. This model solves the difficult cases of face recognition in an image by using the context generated from the caption associated with the same image. Finally, we present an effective solution for classifying the scene shown in an image, which provides useful context for both of the face detection and recognition problems.

George Konidaris; *Autonomous Robot Skill Acquisition;* (Andrew Barto, Advisor); May 2011; Postdoctoral Associate, MIT.

Among the most impressive aspects of human intelligence is skill acquisition–the ability to identify important behavioral components, retain them as skills, refine them through practice, and apply them in new task contexts. Skill acquisition underlies both our ability to choose to spend time and effort to specialize at particular tasks, and our ability to collect and exploit previous experience to become able to solve harder and harder problems over time with less and less cognitive effort. Hierarchical reinforcement learning provides a theoretical basis for skill acquisition, including principled methods for learning new skills and deploying them during problem solving. However, existing work focuses largely on small discrete problems. This thesis identifies the primary obstacles to achieving autonomous skill acquisition in high-dimensional, continuous domains and introduces three methods for overcoming these obstacles: skill chaining, a general skill discovery method for continuous reinforcement learning domains; abstraction selection, an efficient algorithm for selecting a suitable abstraction from a library of available abstractions when creating a new skill; and CST, a method for rapidly building trees of skills (with appropriate abstractions) from sample trajectories obtained via human demonstration, a feedback controller, or a planner. These algorithms are applied to achieve autonomous skill acquisition on the uBot-5.

Ming Li; *Data Management and Wireless Transport for Diverse Sensor Networks;* (Deepak Ganesan and Arun Venkataramani, Advisors); Sept. 2010; Research Staff Member, IBM T.J. Watson.

Today, various sensor networks have emerged and span a wide range of sensing capabilities, computation, energy and communication resources, and user needs. They pose unique design challenges to the distributed system design. We focus on following four challenges in data management and wireless transport. 1) We examine how to explore the resource-rich proxies on the edge of the network to assist the resource-poor sensors. We propose a novel two-tier sensor data management architecture, PRESTO, that proxies model sensed data and predict future data, while sensors check sensed data with model-predicted values. 2) We look at the sensing application regime where a single sensor network has to support diverse users and applications with limited bandwidth and computation resources. We propose a utility-driven approach, MUDS, that maximizes data sharing across users while using limited resources. 3) We seek to improve the matching performance between users' interest and sensed data in large scale sensing applications. We propose BlueDove, a cloud-based publish/subscribe system that takes advantage of the rich resources and flexibility of a computation cloud. 4) We examine how to make the underlying wireless transport between sensor nodes more reliable and efficient. We propose a clean-slate re-design of the network stack, Hop, that uses reliable per-hop block transfer as a building block and builds all other components such as back-pressure congestion control and end-to-end virtual retransmission on top of block transfer.

Daniel Menasche; *Enabling Peer-To-Peer Swarming For Multi-Commodity Dissemination;* (Donald Towsley, Advisor); May 2011; Assistant Professor, Dept. of Computer Science, Federal University of Rio de Janeiro.

Peer-to-peer swarming, as used by BitTorrent, is one of the de facto solutions for content dissemination in today's Internet. Although peer-to-peer swarming has been widely studied for a decade, prior work has focused on the dissemination of one commodity (single file). This thesis focuses on the multi-commodity case. We have discovered through measurements that a vast number of publishers currently disseminate multiple files in a single swarm (bundle). The first contribution of this thesis is a model for content availability. We use the model to show that, when publishers are intermittent, bundling K files increases content availability exponentially as function of K. When there is a stable publisher, we consider content availability among peers. Our second contribution is the estimate of the dependency of peers on the stable publisher (self-sustainability). Then, we investigate reciprocity and the use of barter that occurs among peers. As our third contribution, we prove that the loss of efficiency due to the download of unrequested content to enforce direct reciprocity, as opposed to indirect reciprocity, is at most two in a class of networks without relays. As our fourth contribution, we present conditions for the existence and uniqueness of an equilibrium between publishers and peers.

Albert (Gene) Novark; *Hardening Software Against Memory Errors and Attacks;* (Emery Berger, Advisor); Feb. 2011; Associate, Morgan Stanley.

Programs written in C and C++ are susceptible to a number of memory errors, including buffer overflows and dangling pointers. At best, these errors cause crashes or performance degradation. At worst, they enable security vulnerabilities, allowing denial-of-service or remote code execution. Existing runtime systems provide little protection against these errors. They allow minor errors to crash the program and ensure predictability that allows attackers to consistently exploit vulnerabilities. In this thesis, we introduce a series of runtime systems that detect and tolerate these errors in deployed applications. By design, these systems tolerate minor errors while lowering the probability of successfully exploiting security vulnerabilities. The first such system, Archipelago, protects exceptionally sensitive server applications against severe errors using an object-per-page randomized allocator. It provides near-100% protection against certain common attack vectors. Our second system, DieHarder, combines ideas from Archipelago and previous systems to enable maximal protection against attacks while incurring minimal runtime and memory overhead. Our final system, Exterminator, automatically corrects heap-based buffer overflows and dangling pointers without requiring programmer intervention. Exterminator relies on a low-overhead randomized allocator and statistical inference techniques to isolate and correct errors in deployed applications, deterministically tolerating errors and forcing attackers to adapt to a changing attack surface.

Marek Petrik; *Optimization-based Approximate Dynamic Programming: Reliable Algorithms and Feature Selection;* (Shlomo Zilberstein, Advisor); Sept. 2010; Postoctoral Researcher, IBM Research.

Reinforcement learning algorithms hold promise in many complex domains, such as resource management and planning under uncertainty. Most reinforcement learning algorithms are iterative--they successively approximate the solution based on a set of samples and features. Although these iterative algorithms can achieve impressive results in some domains, they are not sufficiently reliable for wide applicability. Some of the most interesting reinforcement learning algorithms are based on approximate dynamic programming (ADP). This thesis presents new reliable algorithms for ADP that use optimization instead of iterative improvement. We improve on approximate linear programming–an existing method–and derive approximate bilinear programming–a new robust approximate method. The methods presented in this thesis can potentially be used in many practical applications in artificial intelligence, operations research, and engineering. Our experimental results show that optimization-based methods may perform well on resource-management problems and standard benchmark problems and therefore represent an attractive alternative to traditional iterative methods.

Piyanuch (Pla) Silapachote; *Biologically Inspired Vision Systems;* (Allen Hanson, Advisor) May 2011; Faculty, Dept. of Information and Communication Technology, Mahidol University.

One approach to the design of intelligent machines capable of perceiving visual information is to model the fascinating primate vision. Based on discoveries in neuroscience, physiology, and psychology, biologically-plausible models for object recognition and classification can be simple yet achieve high accuracy and generalize well. The proposed system employs unsupervised feature learning, simulating hypercolumns of the primary cortex, a hierarchical feed-forward framework, mimicking simple and complex cells, and neural network classification, a computational model of interconnected neurons. Compared to existing approaches, this system is more biologically inclined as well as more effective. Compared to other state-of-the-art systems it achieves good accuracies with significantly shorter runtimes. Other key features are good generalizability and independence of delicate segmentation procedures employed by many other systems. Experiments are conducted both on natural scenes and challenging realistic underwater marine images. The latter is a rather uncommon data source, no biologically inspired vision systems had previously been applied to it. Despite domain-specific difficulties, such as low image quality, and high diversity of shapes and motions, the potential of the proposed system is shown to be quite promising. This bodes well for future application to other domains currently not considered by mainstream computer vision.

Jacob Sorber; *System Support for Perpetual Mobile Tracking;* (Mark Corner, Advisor); Sept. 2010; Postdoctoral Associate, Dartmouth College.

Advances in low-power electronics, energy harvesting, and sensor technologies will revolutionize mobile and embedded computing, by enabling networks of mobile sensor devices that are long-lived and self-managing. When realized, this new generation of perpetual systems will have a far-reaching and transformative impact, improving scientists' ability to observe natural phenomena, and enabling many ubiquitous computing applications for which regular maintenance is not feasible. Perpetual systems face many challenges. Conditions at runtime are unknown and highly variable. Variations in harvested energy and energy consumption, as well as changes in network connectivity and bandwidth require systems that are able to adapt gracefully at run-time to meet different circumstances. However, when programmers muddle adaptation details with application logic, the resulting code is often difficult to understand as well as maintain. This dissertation demonstrates that perpetual systems can be designed and deployed without sacrificing programming simplicity. We describe two systems. Eon, the first energy-aware language, allows programmers to simply express energy policies and delegate the complexities of energy management to the underlying system. The second, Tula, automatically balances the inherently dependent activities of data collection and data delivery, while also ensuring that devices have fair access to network resources.

Siddharth Srivastava; *Foundations and Applications of Generalized Planning;* (Neil Immerman and Shlomo Zilberstein, Advisors); Sept. 2010; Postdoctoral Research Associate, Dept. of Computer Science, UMass Amherst.

Research in AI Planning is largely focused on the problem of finding plans or sequences of actions for going from a specific initial state to a goal state. The inherent complexity of this task and uncertainty in real-world situations make it desirable to find "generalized" plans which can be used to solve classes of similar problem instances. However, such generalized solutions typically require loops of actions, which are impossible even to verify as correct in the general case because of the undecidability of the halting problem for Turing Machines. In my thesis, I addressed both theoretical and applied aspects of this problem. I first develop an efficient method for determining the correctness and utility of a general class of loops of actions. Next, I utilize this method to develop original algorithms for computing generalized plans by identifying loops in sample classical plans and merging classical plans together; by using classical planners to automate this process and thereby solve a given generalized problem from scratch; and also by conducting a direct search in the space of abstract states.

John Sweeney; *A Teleological Approach to Robot Programming by Demonstration;* (Roderic Gupen, Advisor); Feb. 2011; Research Associate, FDO Partners, LLC.

This dissertation presents an approach to robot programming by demonstration based on two key concepts: demonstrator intent is the most meaningful signal that the robot can observe, and the robot should have a basic level of behavioral competency from which to interpret observed actions. I argue that programming by demonstration can be organized into declarative and procedural components. The declarative component represents a reusable outline of underlying behavior that can be applied to many different contexts. The robot can use the knowledge encapsulated in sensorimotor schemas to interpret the demonstration. The procedural component represents the dynamic portion of the task that is based on features observed at run time. I describe how statistical models, Bayesian methods in particular, can be used to model these components. These models have many features that are beneficial for learning in this domain, such as tolerance for uncertainty, and the ability to incorporate prior knowledge into inferences. I demonstrate this architecture through experiments on a bimanual humanoid robot using tasks from the pick and place domain. Additionally, I develop and experimentally validate a model for generating grasp candidates using visual features that is learned from demonstration data.

Louis Theran; *Problems in Generic Combinatorial Rigidity: Sparsity, Sliders, and Emergence of Components;* (Ileana Streinu, Advisor); Sept. 2010; Research Assistant Professor, Mathematics Dept., Temple University.

A bar-joint framework is made of fixed-length bars connected by universal joints with full rotational degrees of freedom; the allowed motions preserve the lengths and connectivity of the bars, and a framework is rigid if the only allowed motions are trivial motions of Euclidean space. The remarkable Maxwell-Laman Theorem says that rigidity of generic bar-joint frameworks depends only on the graph that has as its edges the bars and as its vertices the joints. We generalize the "degree of freedom counts" that appear in the Maxwell-Laman theorem to the very general setting of (k,l)-sparse and (k,l)-graded sparse hypergraphs, giving graph-theoretic, matroidal, and algorithmic characterization of them. We then introduce a new rigidity model: slider-pinning rigidity. This is an elaboration of the planar bar-joint model to include sliders, which constrain a vertex to move on a specific line. We prove the analogue of the Maxwell-Laman Theorem for slider pinning, using, as a lemma, a new proof of Whiteley's Parallel Redrawing Theorem. We then study the emergence of rigid substructures in a generic framework with the combinatorics of a sparse Erdös-Renyi random graph, proving the existence of a sharp threshold for them to exist and a linear-sized lower bound when they emerge.

Chang Wang; *A Geometric Framework for Transfer Learning Using Manifold Alignment;* (Sridhar Mahadevan, Advisor); Sept. 2010; Research Scientist, IBM TJ Watson Research.

Many machine learning problems involve dealing with a large amount of high-dimensional data across diverse domains. Annotating or labeling the data is also expensive as it involves significant human effort. This work explores a solution to these problems by exploiting the property that high-dimensional data in real-world application domains often lies on a lower-dimensional structure, whose geometry can be modeled as a graph or manifold. We propose a set of novel manifold-alignment based approaches for transfer learning. The proposed approaches transfer knowledge across different domains by finding low-dimensional embeddings of the datasets to a common latent space, which simultaneously match corresponding instances while preserving local or global geometry of each input dataset. We develop several extensions of a manifold alignment framework to more challenging situations, including (1) when no correspondences across domains are given; (2) when the global geometry of each input domain needs to be respected; (3) when label information rather than correspondence information is available. Another contribution of this thesis is the study of multiscale methods for manifold alignment. Multiscale alignment automatically generates alignment results at different levels by discovering the shared intrinsic multilevel structures of the given datasets, providing a common representation across all input datasets.

Philipp Weis; *Expressiveness and succinctness of first-order logic on finite words;* (Neil Immerman, Advisor); May 2011; Software Engineer, Google, Inc.

Expressiveness, and more recently succinctness, are two central concerns of finite model theory and descriptive complexity theory. We develop new bounds on the expressiveness and succinctness of first-order logic with two variables on finite words, present a related result about the complexity of the satisfiability problem for this logic, and explore a new approach to the generalized star-height problem from the perspective of logical expressiveness. Using our complete characterization of the expressive power of first-order logic with two variables on finite words, we prove that the quantifier alternation hierarchy for this logic is strict, settling the main remaining open question about the expressiveness of this logic. We also prove a polynomial-size small-model property for this logic, leading to an NP algorithm and thus proving that the satisfiability problem for this logic is NP-complete. Finally, we investigate the generalized star-height problem. As of today, we do not even know whether there exists a regular language that has generalized star-height larger than 1. We show that this problem can be phrased as an expressiveness question for first-order logic with a restricted transitive closure operator, and use established tools from finite model theory to gain new insights into the generalized star-height problem.