Computer Science
UNIVERSITY OF TORONTO

RESEARCH IN ACTION

Research Highlights

CONTACT US

For more information on
any of these projects:

Research & Industrial Relations Office
Department of Computer Science
10 King's College Road, Room 3303
Toronto, Ontario M5S 3G4
Telephone: 416-978-7390
E-mail: acrir@cs.toronto.edu
Website: www.cs.toronto.edu

In partnership with

NSERC CRSNG

BlackBerry

ARTIFICIAL INTELLIGENCE

COMPUTATIONAL BIOLOGY

COMPUTER GRAPHICS

COMPUTER SYSTEMS AND NETWORKS

DATABASES

HUMAN-COMPUTER INTERACTION

NUMERICAL ANALYSIS

SOFTWARE ENGINEERING

THEORY OF COMPUTATION

www.cs.toronto.edu

MESSAGE FROM THE CHAIR

The Department of Computer Science (DCS) at the University of Toronto, the first CS department in Canada, was established in 1964. Our faculty and students engage in ground-breaking research, teaching, and entrepreneurship that has earned DCS a reputation as one of the top ten Computer Science departments in the world*. Like the field itself, the impact of DCS-incubated research is widespread; with strong interdiscplinary partnerships with the top institutions and researchers in the world, our work extends across areas from finance, to biotechnology, education, business, entertainment, and beyond.

We still believe that, for our department, industrial relations are vital to maintaining our strength and impact as a leader in the field. In order to maintain the strength and impact of our research program, we need to reinforce and cultivate our ties with industry. For that reason, we created our Research & Industrial Relations (RIR) Office five years ago. The Office is currently led by Professor Gerald Penn, and provides the critical liaison between companies and our faculty. We invite you to contact the RIR Office at acrir@cs.toronto.edu.

By getting in touch with the RIR Office, industry can engage with the department in a number of exciting ways: 1) be introduced to researchers working in their area of interest; 2) become acquainted with the growing professional masters program and potentially host an 8-month internship; 3) recruit the best and brightest students; 4) discover other mutually beneficial ways to collaborate with the department and its faculty and students.

The Research in Action showcase also continues to be a great opportunity to establish and bolster these relationships, giving guests a sense of the limitless potential for collaboration across our 11 research groups, not to mention our ongoing work with the world-renowned divisions across the University of the Toronto, the hospitals located a block away, and our partners at peer institutions and companies across the world. I would also like to take this opportunity to welcome and thank our external partners who hosted booths at the 2013 Research in Action event, and, in particular, our event sponsors, NSERC and BlackBerry. We appreciate your support and because of you, we are able to offer this annual opportunity to present the breadth of our research.

I invite prospective partners to explore the breadth and depth of the world-class research in DCS and to share with us their own ideas for new collaborations.

Sven Dickinson
PROFESSOR AND CHAIR

*2012 Shanghai Jiao Tong University Academic Rankings of World Universities [computer science category].

AI: COMPUTATIONAL LINGUISTICS AND NATURAL LANGUAGE PROCESSING

1 Acoustic Modeling Using Deep Convolutional Neural Networks

Deep neural networks have shown great successes for many pattern recognition problems. Acoustic modeling is one example of these successes. In this poster we show that applying deep convolutional neural networks for acoustic modeling outperforms all existing acoustic models. Previous work that used CNN for acoustic modeling applied it only across the time domain ignoring the rich information available on the frequency axis. We also link CNNs to other speaker adaptation techniques.

PROJECT TEAM
Abdel-rahman Mohamed, Graduate Student
Gerald Penn, Faculty

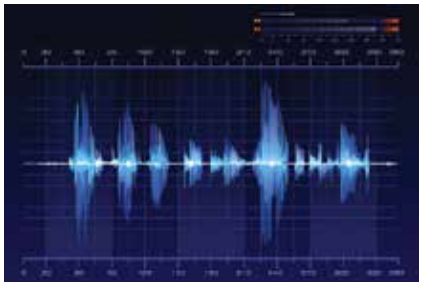

2 Automated Classification of Primary Progressive Aphasia Subtypes from Narrative Speech Samples

In the early stages of neurodegenerative disorders, individuals may exhibit a decline in language abilities that is difficult to quantify with traditional neuropsychological tests. Careful analysis of connected speech can provide valuable information about a patient's language capacities. However, this type of analysis has been limited in the past by its time-consuming nature. In this study, we present a method for evaluating and classifying connected speech in primary progressive aphasia (PPA) using computational techniques. Syntactic and semantic features are automatically extracted from transcriptions of narrative speech for three groups: Semantic Dementia (SD), Progressive Nonfluent Aphasia (PNFA), and healthy controls. Features that vary significantly between the groups are used to train machine learning classifiers, which are then tested on held-out data. We achieve accuracies above baseline on the three classification tasks. An analysis of the influential features shows that the computationally extracted measures agree well with previous findings in the PPA literature, suggesting that this method could have valuable applications to diagnosis and research.

PROJECT TEAM
Katie Fraser, Graduate Student
Graeme Hirst, Faculty
Jed Meltzer, Rotman Research Institute
Naida Graham, Toronto Rehabilitation Institute
Carol Leonard, University of Ottawa
Sandra Black, Sunnybrook Health Sciences Centre
Elizabeth Rochon, Dept. of Speech-Language Pathology

3 Distinguishing Automatic from Human-Written Summaries by Information Content

Automatic summarization is currently dominated by extractive systems that concatenate snippets of source text to form a summary, yet the gold-standard human-written summaries against which they are evaluated are abstractive, requiring domain knowledge and semantic inference to decide what information is important and how this should be expressed in the summary. We attempt to quantify the extent to which extractive systems can approximate this semantic inference by comparing dependency-based semantic representations of basically extractive, state-of-the-art summaries to abstractive model summaries. Our analysis suggests that current systems have over-fit to the topical dependency edges found in the source text, and that a better way to emulate model summaries is to consider sentence aggregation approaches using domain knowledge.

PROJECT TEAM
Jackie Chi Kit Cheung, Graduate Student
Gerald Penn, Faculty

AI: COMPUTATIONAL VISION

4 Discovering Categorical Shape from Captioned Images

Visual object categorization systems typically require a strong degree of supervision for model learning, relying on bounding boxes drawn around the object of interest to crop a cluttered scene. Captioned images are ubiquitous and offer a semi-supervised approach for learning category models without the need for prior annotation. To address this challenging task we apply a recent language-vision integration framework that finds spatial configurations of image features that co-occur with words in image captions. By using local invariant contour features, objects are recognized by a dynamic model that grows along the object's boundary. Experiments on the ETHZ shape dataset demonstrate 1) an improvement by using contours to capture object categories, and that 2) a typical approach to category recognition relies heavily on bounding box annotation.

PROJECT TEAM
Tom Sie Ho Lee, Graduate Student
Sanja Fidler, Postdoctoral Fellow
Alex Levinshtein, Postdoctoral Fellow
Sven Dickinson, Faculty

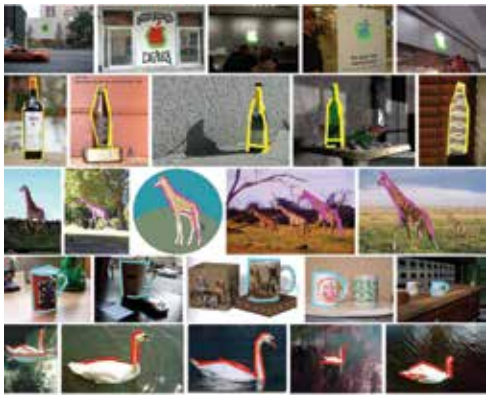

www.cs.toronto.edu

5 Efficient Coordinated Power Distribution on Private Infrastructure

Due to the design of current power distribution networks, it will become increasingly attractive for agents to generate their own power (distributed generation) and to construct private infrastructure (e.g., transmission lines) to exchange power with others nearby without using the main public grid. We show that such private transactions may increase overall load on the network because of the increased distance that power must flow from generation sources, thus increasing transmission loss. We present a coordination scheme that allows centralized control of private infrastructure while satisfying participation constraints and budget balance. Experiments show that our scheme reduces transmission losses by 4-5% when there are only a constant number of private lines and by 55%-60% when the number of private lines is proportional to the number of agents. These results hold even when there is only a small amount of distributed generation

PROJECT TEAM
Andrew Perrault, Graduate Student
Craig Boutilier, Faculty

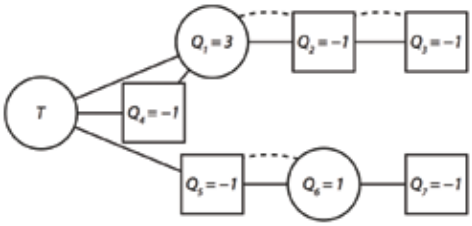

6 Multi-Dimensional Single-Peakedness and its Approximations

Single-peakedness is the most common domain restriction in social choice. However, the extent to which agents' preferences are single-peaked in practice, and whether recent proposals for approximate single-peakedness can help, is unclear. We assess the ability of both single-dimensional and multi-dimensional approximations to explain preference profiles drawn from several real world elections. We develop a simple branch-and-bound algorithm that finds multi-dimensional, single-peaked axes that best fit a given profile, and which works with several forms of approximation. Empirical results on two election data sets show that preferences in these elections are far from single-peaked in any one-dimensional space, but are very nearly single-peaked in two dimensions. Our algorithms are reasonably efficient, and also show excellent anytime performance.

PROJECT TEAM
Xin Sui, Graduate Student
Alex Francois-Nienaber, Undergraduate student
Craig Boutilier, Faculty

7 Robust Matching with Partial Preferences: Elicitation and Approximately Stable Matching

Algorithms for stable marriage and related matching problems typically assume that full preference information is available, or have restrictive assumptions regarding incomplete information. While the Gale-Shapley algorithm can be viewed as a means of eliciting preferences incrementally, it does not prescribe a general means for matching with incomplete information, nor is it designed to minimize elicitation. We propose the use of maximum regret to measure the degree of instability of a matching with partial preferences; minimax regret to find matchings that are maximally stable in the presence of partial preferences; and heuristic elicitation schemes that use max regret to determine relevant preference queries. Several of the schemes presented find stable matchings while eliciting considerably less preference information than Gale-Shapley and are much more appropriate in settings where approximate stability is viable, as these schemes converge quite rapidly.

PROJECT TEAM
Joanna Drummond, Graduate Student
Craig Boutilier, Faculty

AI: MACHINE LEARNING AND KNOWLEDGE REPRESENTATION

8 Empathetic Social Choice on Social Networks

We introduce a model for social choice, specifically consensus decision making on social networks that reflects dependence among the utilities of connected agents. We define an empathetic social choice framework in which agents derive utility based on both their own intrinsic preferences and the satisfaction of their neighbors. We show how this problem translates into a weighted form of classical preference aggregation (e.g. social welfare maximization or voting) and develop scalable algorithms for consensus decision making. Empirical results validate the effectiveness of our methods and demonstrate the value of accounting for empathetic preferences.

PROJECT TEAM
Amirali Salehi-Abari, Graduate Student
Craig Boutilier, Faculty

9 Image Segmentation with Compositional High Order Constraints

In this work, we study the learning of a general class of pattern-like high order potential, which we call Compositional High Order Pattern Potentials (CHOPPs). We show that CHOPPs include the linear deviation pattern potentials of Rother et al. [23] and also Restricted Boltzmann Machines (RBMs); we also establish the near equivalence of these two models. We then improve CHOPPs performance in high variability data sets with two primary contributions: (a) developing a loss-sensitive joint learning procedure, so that internal pattern parameters can be learned in conjunction with other model potentials to minimize expected loss and (b) learning an image-dependent mapping that encourages or inhibits patterns depending on image features.

PROJECT TEAM
Yujia Li, Graduate Student
Richard Zemel, Faculty

www.cs.toronto.edu

COMPUTATIONAL BIOLOGY

10 A Novel Structurally Regularized Multi-Task Lasso Model for Phenotype Prediction

With the recent advances in DNA sequencing technologies, we can now obtain the genomic profile of an individual easily. The next big challenge is making useful predictions from this data. In particular, predicting one's phenotype, such as disease risk predisposition, height and weight, solely based on his/her genome is one of the biggest promises of personalized genomics. However, this so called phenotype prediction problem proved to be a difficult one due to immense complexity in the underlying genetics. Recently, several machine learning algorithms for phenotype prediction have been proposed. In particular, the observation that several complex traits can be better explained by additive (small) effects of several genetic loci led to a community-wide interest in designing sparse linear prediction models, primarily Lasso models. In this project, we propose a novel multi-task Lasso model that leverages omics data for regularization. Our model makes use of dependencies in both the input and the output space in the form of structural regularizes. Our integrative Lasso model obtains biologically more meaningful and easily interpretable set of features for improved prediction.

PROJECT TEAM
Recep Colak, Graduate Student
Philip Kim, Faculty
Karsten Borgwardt, Max Planck Institute, Tuebingen, Germany
Barbara Rakitsch, Max Planck Institute, Tuebingen, Germany

COMPUTER GRAPHICS

11 CrossShade

CrossShade allows the creation of 3D-looking shaded production drawings from concept sketches. We use artist-drawn cross-section lines to automatically infer surface information across the sketch, enabling 3D-like rendering. Cross-sections function as an aid to both sketch creation and viewer understanding of the depicted 3D shape. In particular, intersections of these curves, or cross-hairs, convey valuable 3D information that viewers compose into a mental model of the sketch. We use this information to estimate the surface normals and render shaded images.

PROJECT TEAM
Cloud Shao, Graduate Student
Adrien Bousseau, Researcher, INRIA Sophia Antipolis
Alla Sheffer, Faculty, UBC
Karan Singh, Faculty

COMPUTER SYSTEMS AND NETWORKS

12 Accelerating the Cloud with Heterogeneous Computing

Heterogeneous multiprocessors that combine multiple CPUs and GPUs on a single die are becoming commonplace in the market. As seen recently from the high performance computing community, leveraging a GPU can yield performance increases of several orders of magnitude. We propose using GPU acceleration to greatly speed up cloud management tasks in the cloud operating system - The Virtual Machine Monitor. This is only becoming possible now that the GPU is moving on - chip, since the latency across the PCIe bus was too great to make fast, informed decisions about the state of a system at any given point. We explore various examples of cloud management tasks that can greatly benefit from GPU acceleration. We also tackle tough questions of how to manage this hardware in a multi-tenant system. Our proof-of-concept case study shows that for a common cloud operation, memory deduplication, GPU acceleration can improve the performance of its hashing component by a factor of over 80.

PROJECT TEAM
Sahil Suneja, Graduate Student
Eyal de Lara, Faculty

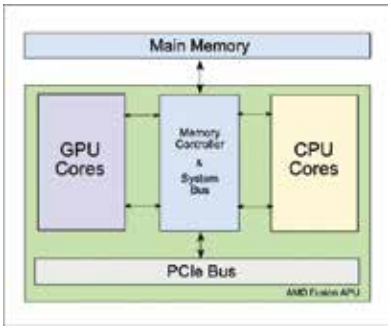

13 Bursty Subgraphs in Social Networks

Data available through social media and content sharing platforms, present opportunities for analysis and mining. In the context of social networks, it is interesting to formalize and locate bursts of activities amongst users, related to a particular event and to report sets of socially connected users participating in such bursts. Such collections present new opportunities for understanding social events, and render new ways of online marketing. In this work, we model social information using two conceptualized graph models. The first one (the action graph) provides a detailed model of all activities of all users while the second one (the holistic graph) provides an aggregate view on each user in the social media. We also propose two models to define the notion of "burst". The first model (intrinsic burst model) takes the intrinsic characteristics of each user into account to recognize the bursty behaviors; while the second model (social burst model) considers neighbours' influences when identifying bursts. We provide two linear algorithms to detect bursts based on the proposed models. These algorithms have been extensively evaluated on a month of full Twitter dataset certifying the practicality of our approach. A detailed qualitative study of our techniques will also be presented.

PROJECT TEAM
Milad Eftekhar, Graduate student
Yashar Ganjali, Faculty

www.cs.toronto.edu

14 Smarter Home Networks with Simple Gateways

Users demand smarter home networks that conveniently provide networking features. Previous attempts focus on pushing more features into the gateways and making them smarter. Home gateways, however, have limited resources and cannot accommodate resource-demanding features. In this project, in contrast, we delegate most gateway functionalities to a shared, consolidated control platform, and keep the gateway as simple as possible. Using this framework, we are not restricted by resource limitations of home gateways and thus are able to provide services such as device discovery or troubleshooting with less effort and better quality.

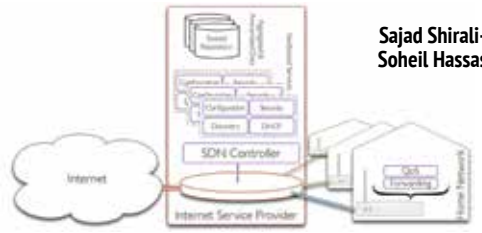

PROJECT TEAM
Sajad Shirali-Shahreza, Graduate Student
Soheil Hassas Yeganeh, Graduate Student
Yashar Ganjali, Faculty

15 Unlocking the Future Wireless Networks

We present a new technology that allows to create high performance wireless networks in a cheaper and simpler way than it is currently possible. The technology can be used to support emerging wireless applications such as medical and health applications, smart grid applications, and smart homes, as well as to increase the capacity of traditional wireless cellular networks through the creation of smaller cells. The key challenges in building these future wireless networks is to design a simple decentralized algorithm for managing the interference between wireless devices that is able to achieve a high performance in terms of bandwidth and delay. We developed in our research group such a mechanism which has the best performance among the currently known protocols.

PROJECT TEAM
Amir Hejazi, Graduate Student
Peter Marbach, Faculty

16 VMM-Based Cloud Analytics with Real-Time Memory Introspection

The adoption of virtualization in modern data centers has brought with it new opportunities and challenges for system monitoring and management. We argue that the dramatic proliferation of virtual machines (VMs) and the highly-dynamic and transient nature of VM provisioning is incompatible with existing monitoring and management solutions that are designed for long-running and highly-available physical systems. Instead, we make the case for promoting the Virtual Machine Monitor (VMM) itself as a universal agent for managing virtualized systems. We describe our memory introspection based approach to manage and monitor VMs from the outside without having any hooks installed inside the VMs. We review the challenges inherent with this approach, the most important being the semantic gap between the VM-internal and the VMM-external view, and describe how we address these challenges in our proposed solution. We discuss new opportunities this approach creates for cluster-scale analytics. We demonstrate the feasibility and potential of our proposed approach with a proof-of-concept prototype of a real-time remote VM-resource monitor.

PROJECT TEAM
Sahil Suneja, Graduate Student
Eyal de Lara, Faculty

17 Exploring Memory Residency for Database Indexes on Variable-Length Keys

In recent years, the size of database memory has increased to the point where it is common for an index to fit into main memory, thus avoiding costly disk accesses. Consequently, indexes must be (and have been) redesigned for memory residency. Yet for indexes with variable-length or long keys, there is a particular tension between the amount of space required to store these keys and the limited space in memory and cache, and there are few indexes designed to efficiently resolve this tension. To guide future research, we present a characterization of a handful of valuable design decisions, and evaluate them both in isolation and in combination with one another.

PROJECT TEAM
Benjamin Reilly, Graduate Student
Zhaohan Daniel Guo, Undergraduate Student
Per-Thomas Vik Parker, Undergraduate Student
Ryan Johnson, Faculty

18 NSERC Business Intelligence Network (BIN)

A five year research network focusing on data analytics/business intelligence (2009-2014). The goal of the NSERC Business Intelligence Network is to create an innovation platform for pre-competitive research in Business Intelligence in Canada. NSERC BIN aims to enhance Canadian business competitiveness through the development of intelligent data management and decision-making solutions. The NSERC BIN research team consists of 15 investigators at seven universities (University of Alberta, University of British Columbia, Carleton University, Dalhousie University, University of Ottawa, University of Toronto, and University of Waterloo). Investigators are working in partnership with researchers at a number of organizations including SAP Canada Inc., IBM Canada Limited, iAnywhere Solutions Canada Ltd., Palomino System Innovations Inc. and Zerofootprint.

PROJECT TEAM
Leslie Dolman, Executive Director, BIN
Annette Mayer, Network Manager, BIN
Renée Miller, Faculty

NOTES

www.cs.toronto.edu

19 Conductor: Managing Applications in a Multi Device Environment

Being surrounded by computing devices, in their many forms, is rapidly fading from being an intriguing novelty to a familiar part of our daily lives. It is not uncommon for a person to make use of a significant number of devices such as smart phones, tablets, and televisions. The structure of interactive software is shifting towards interactions that are not device specific. Applications that revolve around the user,where each device is less a container for applications and more an instrument through which applications can be seen and heard. For instance, interfaces for cross device applications should be flexible and easily separated between devices. A user might use a tablet to select the pictures being displayed on a larger display. We present Conductor, an interface for managing and establishing cross device applications. Conductor does not depend on real world positions to create links between devices. Conductor forms a network of devices through device sensors such as accelerometers, WIFI, and camera sensors. Conductor also introduces a number of interaction techniques for passing information between devices and controlling devices. The system also visualizes application sessions, allowing a user to easily put down one device and continue working on the same session with a different device. Two cross device applications were built to evaluate Conductor, an email application and a map application. These applications highlight the potential benefits of taking advantage of the varied capabilities of different devices.

PROJECT TEAM
Peter Hamilton, Graduate Student
Daniel Wigdor, Faculty

20 Designing Pen Interfaces that Adapt to Device Capabilities

This project presents a tablet application with an interface that is specially tailored for the strengths and weaknesses of pen and touch interaction, while adapting for styluses with varying capabilities. Touch - on its own is well suited to the consumption of content, but cumbersome in its ability to act as an input tool for content creation. There have long existed active-digitizer styluses that exhibit capabilities like pen pressure levels, hover sensing, tilt, buttons and differentiation from touch input. Much research effort has been placed on designing improved interaction that uses these capabilities, often with the goal of offering the pen as a more complete alternative for content creation on multi-touch devices. A new trend is the rising popularity of basic pens for standard touchscreens whose tips simply mimic finger touches, but lack the features offered by the more sophisticated styluses. By presenting a series of principles for designing pen interaction across the spectrum of device capabilities, alongside an example application that implements these patterns, we hope to outline a way forward for content creation interfaces on mobile devices.

PROJECT TEAM
Ben McCanny, Undergraduate Student
Daniel Wigdor, Faculty

21 Interactive Exploration of Faceted Information Spaces

The information resources we encounter nowadays are not only large in sizes but also complicated in structures. For example, online academic library databases not only include large number of items, such as publications, but also references (links) across these publications and a range of attributes (facets) on each item, such as authors and keywords. To address the need of efficient browsing and analyzing of such dataset, we present an interactive visualization system for faceted information spaces, which supports progressive visual exploration of the data by allowing users to dynamically query, filter, manipulate and correlate the data items based on their attributes and interlinks.

PROJECT TEAM
Jian Zhao, Graduate student
Ravin Balakrishnan, Faculty

22 Panelrama: A Distributed UI Framework for Multiple Devices

With the emergence of new computing form factors like tablets and smartphones, it is increasingly likely for consumers to own two or more computing devices. Currently the majority of applications on these devices operate independently, making it difficult for users to take advantage of all their devices at the same time. The possibility of sharing application state and I/O resources between devices opens up a wealth of interaction possibilities: a user may retain workflow between devices, annex additional hardware to overcome device limitations, or use additional devices to enrich interactions.

Our solution is Panelrama, a web framework which allows users to interact with an application across multiple devices simultaneously. Panelrama distributes subsections of the complete application interface to each device, allowing individual devices to take on functions ideal for their role. For example, a video streaming application may allocate its video stream to a TV, playback controls to a smartphone, and related videos to a tablet. We will present prototype applications built using the Panelrama framework, including distributed satellite map and video streaming applications.

PROJECT TEAM
Jishuo Yang, Graduate Student
Daniel Wigdor, Faculty

www.cs.toronto.edu

23 Designing Novel Interaction Methods for Mobile Data Input and Tracking

Every time we interact with technology we leave scraps of information behind: Where we've been, whom we're talking to, what we're doing, how we're feeling and more. Our personal devices can now know and remember more about our day than we do ourselves. In the project being proposed, we will study how these digital information scraps can be used to improve self-reflection and promote healthy habits. Many papers have noted the benefits of self-tracking metrics in your life such as your sleep, time of eating, mood and physical activity. However, existing tracking methods are cumbersome and interrupt your daily flow. This project will create and study a novel, low-friction method of facilitating self-tracking by allowing the user to input data while unlocking their phone and other devices. Daily summaries will allow users to see how they have progressed in their goals and tracking.

PROJECT TEAM
Thariq Shihipar, Undergraduate Student, Engineering
Daniel Wigdor, Faculty

24 TAGLab - Technologies for Aging Gracefully Lab

TAGlab is comprised of talented individuals with backgrounds in computer science, engineering, human-computer interaction, graphic and interface design, and psychology. We work with researchers and clinicians to find ways that digital media can help people remain vigorous and independent, strengthen ties to family and community, and preserve their identity as they age. Our mission is R&D in support of aging throughout the life course. We identify "sweet spots" where technology seems relevant to human need, envision ways in which the technology could address those needs, then design and test prototypes.

TAG lab

PROJECT TEAM
Abbas Attarwala, Graduate Student
Jessica David, Graduate Student
Carrie Demmans Epp, Graduate Student
Natalie Langlois, Graduate Student
Velian Pandeliev, Graduate Student
Katherine Sellen, Postdoctoral Fellow
Stephen Tsourounis, Undergraduate Research Assistant
Ronald M. Baecker, Faculty

NOTES

25 Touching the Third Dimension

Augmented Reality (AR) is widely employed in interactive systems to seamlessly blend the virtual world with the physical world and enable users to visualize and manipulate virtual objects physically in an intuitive and creative manner. Visualization technologies now allows us to over-impose the physical and virtual world in a accur AR researches have largely focused on enhancing realistic freehand interaction. While these systems allows for increasingly more convincing interactions with virtual objects, they are designed to provide users with only visual feedbacks of the virtual scene, foregoing the haptic feedback that is an integral part of the real-world interaction experience. Users cannot physically "feel" the virtual objects they are manipulating, thus preventing such systems from creating a stronger AR experience. Our prototype is an interactive AR system that employs depth-sensing devices such as Kinect and/or Leap Motion and a programmable robotic arm to provide users with not only visual outputs, but also haptic feedbacks. The computer screen behaves as a fish-tank window into the virtual world that exists behind the monitor; and users can manipulate on-screen objects by reaching behind the monitor into the virtual world. By tracking user's hand positions with high-precision and in real-time, as user's hand approaches the on-screen 3D object, the robotic arm becomes a physical proxy of the virtual object the user is about to touch. In other words, it will move accordingly to user's hand position to provide users with haptic feedback as if the virtual object actually existed behind the monitor. With a commodity robotic arm, this prototype provides a workspace that can easily support a variety of natural haptic feedbacks that range from touching objects, to pressing buttons, lifting/grasping objects, or, with additional hardware, feeling textures.

PROJECT TEAM
Eric Yao, Undergraduate Student
Ricardo Jota, Postdoctoral Fellow
Daniel Wigdor, Faculty

NUMERICAL ANAYLSIS

26 Parallel Implementation on GPUs of an ADI Method for Solving Coupled Parabolic PDEs from Finance

We develop a parallel implementation on a graphics processing unit (GPU) of an alternating dimension implicit (ADI) method for solving coupled linear parabolic partial differential equations in three dimensions with mixed spatial derivatives. We use non-uniform finite differences for the spatial discretization of the PDE and an ADI scheme with stabilizing corrections for time discretization. We apply this method to the problem of pricing European options in three-factor models with stochastic volatility driven by a regime-switching process.

PROJECT TEAM
Kirill Ignatiev, Graduate Student
Duy Minh Dang, Postdoctoral Fellow, Waterloo
Ken Jackson, Faculty

www.cs.toronto.edu

27 A Parallel Non-Linear Iterative Reconstruction Algorithm for X-Ray Computerized Tomography

X-Ray Computerized Tomography (CT) is an imaging technique used to produce tomographic images of various inner parts of a system. It is widely used in medical CT scans as well as various industries. The basic idea is to shoot X-Ray beams at an object from various angles. Depending on the shape and type of the object, these X-Ray beams loose some portion of their initial energy while passing through the object. Based on this change in the energy, one can create algorithms to construct relatively accurate images of the inner tomography of the object. In this project we build up on the work a Nargol Rezvani (a former DCS PhD student) and modify her existing reconstruction algorithm to implement it on parallel CPU's or GPU's using OpenCL and MATLAB. Also, we suggest a slightly different procedural algorithm with better time and space complexity using algorithms similar to the popular Levenberg Marquardt-Fletcher algorithm.

PROJECT TEAM
Mehrdad Shahi, Undergraduate Student
Rinat Abdrashitov, Undergraduate Student
Nargol Rezvani, DCS Alumna
Ken Jackson, Faculty

SOFTWARE ENGINEERING

28 Designing Casual Video Games

The video game community in the GTA is largely composed of small to medium game design studios. One of the biggest audiences for these studios is the casual gaming market, where games are designed to be quick, accessible and social. Designing these games is an exercise in both software engineering and human-computer interaction, to create this appealing experience for a casual audience.

PROJECT TEAM
Osman Haque, Undergraduate Student
Harrison Dahme, DCS Alumnus & Senior Software Engineer, Zynga
Somang Nam, Undergraduate Student
Brent Mombourquette, Undergraduate Student
Abhishek Chaudhry, OCAD U student
Charlotte Gao, OCAD U
Alice Li, DCS Alumna, Ubisoft
Katie Poon, OCAD U student
Steve Engels, Faculty

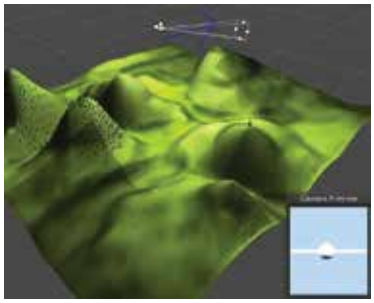

29 Gamifying Collaborative Decision Making

Engaging people to participate fully in an online collaborative decision-making activity can be challenging. Gamification has been used in a variety of online environments to incent and increase participation. In this project, we investigate how gamification can be used to incent and motivate people to participate in and contribute to online collaborative decision-making. We developed two decision-making tools based on thinklets, which are re-usable collaboration activities that facilitate patterns of thinking among people working toward a goal. We integrated our tools into SAP Streamwork, a web-based collaborative platform. We added game dynamics to each of these tools and designed an experiment with a graduate level project management class. We asked 18 groups of students to engage in a project selection decision-making task. Each group randomly used the system with or without gamification. After the experiment we compared quantity and quality of contributions between groups that used gamified system and groups that used the system normally. We show that when game elements are used properly they a) will have significant effect on quantity of contributions in brainstorming tool; and, b) will have significant effect on user engagement in online conversation during the decision-making process.

PROJECT TEAM
Mohammad Ali Moradian, Graduate Student
Maaz Nasir, Graduate Student
Rock Leung, SAP Canada
Kelly Lyons, Faculty

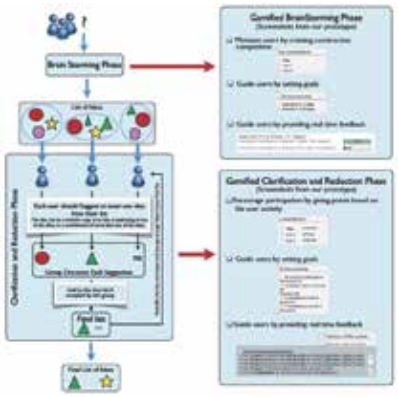

30 Incorporating Service Learning in Introductory Software Engineering

"Introduction to Software Engineering" (CSC301) is a third year course where students learn agile software development techniques and apply them in a software project for a customer. The result is a practical, hands-on course with the benefit of classic software engineering concepts and the vagaries of real-world projects.

PROJECT TEAM
Susan Elliott Sim, Sessional Lecturer, Undergraduate Program

www.cs.toronto.edu

31 MarkUs

MarkUs is an open-source web application that aids in the submission and grading of student assignments. It manages student group formation, allows students to submit their work either through the web interface or through a version control system, and most importantly provides an interface for graders to annotate and grade student work. MarkUs is used by thousands of students in at least three universities in Canada and France. MarkUs has been entirely built by undergraduate students and has proven to be a great training ground for student developers to learn the software engineering and project management skills required when working on a large project with real users.

PROJECT TEAM
Karen Reid, Faculty

32 Using the Model Management Tool Framework for Managing Model Relationships

Model management has emerged as an important approach to addressing the complexity introduced by the use of many interrelated models in Software Engineering. In this tool demonstration, we highlight the capabilities of our Model Management Tool Framework (MMTF) -- an Eclipse-based pluggable infrastructure for rapidly developing model management tools. MMTF differs from other model management frameworks through its emphasis on interactive model management, graphical mega-modeling and strong typing. We illustrate the use of MMTF in the context of automotive software development.

PROJECT TEAM
Michalis Famelis, Graduate Student
Alessio Di Sandro, Research Engineer
Rick Salay, Postdoctoral Fellow
Marsha Chechik, Faculty

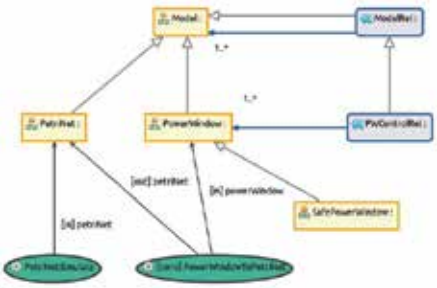

33 Encryption Systems Enabling Fine-Grained Access Control to the Data

Traditional encryption schemes provide a binary access control to the data: a user who holds a secret key to a ciphertext can decrypt and learn the data, while others learn nothing about it. However, in the modern world, where users can play multiple roles and data should be learnable by the "selected few", expressing rich access control policies using these traditional schemes is infeasible. For example, a patient might want to encrypt his personal information and medical history such that each doctor can learn only parts of the data related to his/her specialty. In this work, we present new advances in the Attribute-based Encryption schemes which enable efficient fine-grained access control capabilities and more.

PROJECT TEAM
Sergey Gorbunov, Graduate Student
Vinod Vaikuntanathan, Faculty

34 New Greedy Heuristics for Set Cover and Set Packing

The Set Cover problem (SCP) and Set Packing problem (SPP) are standard NP-hard combinatorial optimization problems. Their decision problem versions are shown to be NP-Complete in Karp's 1972 paper. We specify a rough guide to constructing approximation heuristics that may have widespread application and apply it to devise greedy approximation algorithms for SCP and SPP where the selection heuristic is a variation of that in the standard greedy approximation algorithm. We prove that the technique we use for determining a valuation of the input sets yields a unique value for all Set Cover instance. For both SCP and SPP we give computational evidence that the valuations we specify are unique and can be computed to high precision quickly by an iterative algorithm. Others have experimented with testing the observed approximation ratio of various algorithms over a variety of randomly generated instances, and we have extensive experimental evidence to show the quality of the new algorithm relative to greedy heuristics in common use. Our algorithms are somewhat more computationally intensive than the standard ones, though they are still practical for large instances. We discuss some ways to speed up our algorithms that do not significantly distort their effectiveness in practice on random instances.

PROJECT TEAM
David Kordalewski, Graduate Student
Allan Borodin, Faculty
Ken Jackson, Faculty

NOTES

www.cs.toronto.edu

Abdrashitov, Rinat (Project #27)
UNDERGRADUATE STUDENT
rinat.abdrashitov@mail.utoronto.ca

Attarwala, Abbas (#24)
GRADUATE STUDENT
attarwal@cs.toronto.edu

Baecker, Ronald (#24)
FACULTY
ron@taglab.ca

Balakrishnan, Ravin (#21)
FACULTY
ravin@dgp.toronto.edu

Borodin, Allan (#34)
FACULTY
bor@cs.toronto.edu

Boutilier, Craig (#5, #6, #7, #8)
FACULTY
cebly@cs.toronto.edu

Chaudhry, Abhishek (#28)
OCAD U STUDENT
ac09fq@student.ocadu.ca

Chechik, Marsha (#32)
FACULTY
chechik@cs.toronto.edu

Cheung, Jackie Chi Kit (#3)
GRADUATE STUDENT
jcheung@cs.toronto.edu

Colak, Recep (#10)
GRADUATE STUDENT
rcolak@cs.toronto.edu

Dahme, Harrison (#28)
ALUMNUS
harrison.dahme@gmail.com

Dang, Duy Minh (#26)
POSTDOCTORAL FELLOW, WATERLOO
dm2dang@uwaterloo.ca

David, Jessica (#24)
GRADUATE STUDENT
jdavid@cs.toronto.edu

de Lara, Eyal (#12, #16)
FACULTY
delara@cs.toronto.edu

Demmans Epp, Carrie (#24)
GRADUATE STUDENT
carrie@taglab.ca

Di Sandro, Alessio (#32)
RESEARCH ENGINEER
disandro@cs.toronto.edu

Dickinson, Sven (#4)
FACULTY
sven@cs.toronto.edu

Dolman, Leslie (#18)
EXECUTIVE DIRECTOR, BIN
Leslie.dolman@utoronto.ca

Drummond, Joanna (#7)
GRADUATE STUDENT
jdrummond@cs.toronto.edu

Eftekhar, Milad (#13)
GRADUATE STUDENT
milad@cs.toronto.edu

Engels, Steve (#28)
FACULTY
sengels@cs.toronto.edu

Famelis, Michalis (#32)
GRADUATE STUDENT
famelis@cs.toronto.edu

Fidler, Sanja (#4)
POSTDOCTORAL FELLOW
fidler@cs.toronto.edu

Francois-Nienaber, Alex (#6)
UNDERGRADUATE STUDENT
lex@cs.toronto.edu

Fraser, Katie (#2)
GRADUATE STUDENT
kfraser@cs.toronto.edu

Gangali, Yashar (#13, #14)
FACULTY
yganjali@cs.toronto.edu

Gao, Charlotte (#28)
OCAD U STUDENT
cg08qj@student.ocadu.ca

Gorbunov, Sergey (#33)
GRADUATE STUDENT
sgorbunov@cs.toronto.edu

Guo, Zhaohan Daniel (#17)
UNDERGRADUATE STUDENT
daniel.guo@mail.utoronto.ca

Hamilton, Peter (#19)
GRADUATE STUDENT
peter.hamilton@mail.utoronto.ca

Haque, Osman (#28)
UNDERGRADUATE STUDENT
osmanaminhaque@gmail.com

Hassas Yeganeh, Soheil (#14)
GRADUATE STUDENT
soheil@cs.toronto.edu

Hejazi, Amir (#15)
GRADUATE STUDENT
amir@cs.toronto.edu

Hirst, Graeme (#2)
FACULTY
gh@cs.toronto.edu

Ignatiev, Kirill (#26)
GRADUATE STUDENT
ikirill@cs.toronto.edu

Jackson, Ken (#26, #27, #34)
FACULTY
krj@cs.toronto.edu

Johnson, Ryan (#17)
FACULTY
ryan.johnson@cs.utoronto.ca

Jota, Ricardo (#25)
POSTDOCTORAL FELLOW
jotacosta@dgp.toronto.edu

Kim, Philip (#10)
FACULTY
pm.kim@utoronto.ca

Kordalewski, David (#34)
GRADUATE STUDENT
davekordalewski@gmail.com

Langlois, Natalie (#24)
GRADUATE STUDENT/ADMINISTRATOR
natalie@taglab.ca

Lee, Tom Sie Ho (#4)
GRADUATE STUDENT
tshlee@cs.toronto.edu

Leung, Rock (#29)
SAP CANADA
rockleung @ gmail.com

Levinshtein, Alex (#4)
POSTDOCTORAL FELLOW
babalex@cs.toronto.edu

Li, Alice (#28)
ALUMNA
li.alicem@gmail.com

Li, Yujia (#9)
GRADUATE STUDENT
yujiali@cs.toronto.edu

Lyons, Kelly (#29)
FACULTY
kelly.lyons@utoronto.ca

Marbach, Peter (#15)
FACULTY
marbach@cs.toronto.edu

Mayer, Annette (#18)
NETWORK MANAGE,R BIN
annette.mayer@utoronto.ca

McCanny, Ben (#20)
UNDERGRADUATE STUDENT
ben.mccanny@mail.utoronto.ca

Miller, Renée (#18)
FACULTY
miller@cs.toronto.edu

Mohamed, Abdel-rahman (#1)
GRADUATE STUDENT
asamir@cs.toronto.edu

Mombourquette, Brent (#28)
UNDERGRADUATE STUDENT
g0mombou@cdf.toronto.edu

Moradian, Mohammad Ali (#29)
GRADUATE STUDENT
moradian@cs.toronto.edu

Nam, Somang (#28)
UNDERGRADUATE STUDENT
g0somang@cdf.toronto.edu

Nasir, Maaz (#29)
GRADUATE STUDENT
maaz.nasir@utoronto.ca

Pandeliev, Velian (#24)
GRADUATE STUDENT
vpandeli@cs.toronto.edu

Penn, Gerald (#1, #3)
FACULTY
gpenn@cs.toronto.edu

Perrault, Andrew (#5)
GRADUATE STUDENT
perrault@cs.toronto.edu

Poon, Katie (#28)
OCAD U STUDENT
kp09pe@student.ocadu.ca

Reid, Karen (#31)
FACULTY
reid@cs.toronto.edu

Reilly, Benjamin (#17)
GRADUATE STUDENT
breilly@cs.toronto.edu

Rezvani, Nargol (#27)
ALUMNA
nrezvani@cs.toronto.edu

Salay, Rick (#32)
POSTDOCTORAL FELLOW
rsalay@cs.toronto.edu

Salehi-Abari, Amirali (#8)
GRADUATE STUDENT
abari@cs.toronto.edu

Sellen, Katherin (#24)
POSTDOCTORAL FELLOW, ENGINEERING
kate@taglab.ca

Shahi, Mehrdad (#27)
UNDERGRADUATE STUDENT
mehrdad.shahi@mail.utoronto.ca

Shao, Cloud (#11)
GRADUATE STUDENT
riashowcase@cloudshao.com

Shihipar, Thariq (#23)
UNDERGRADUATE STUDENT, ENGINEERING
thariq.shihipar@mail.utoronto.ca

Shirali-Shahreza, Sajad (#14)
GRADUATE STUDENT
shirali@cs.toronto.edu

Sim, Susan Elliott (#30)
SESSIONAL LECTURER
ses@drsusansim.org

Singh, Karan (#11)
FACULTY
karan@dgp.toronto.edu

Sui, Xin (#6)
GRADUATE STUDENT
xsui@cs.toronto.edu

Suneja, Sahil (#12 #16)
GRADUATE STUDENT
sahil@cs.toronto.edu

Tsourounis, Stephen (#24)
UNDERGRADUATE RESEARCH ASSISTANT
steve@taglab.ca

Vaikuntanathan, Vinod (#33)
FACULTY
vinodv@cs.toronto.edu

Vik Parker, Per-Thomas (#17)
UNDERGRADUATE STUDENT
per.parker@utoronto.ca

Wigdor, Daniel (#19,#20, #22, #23, #25)
FACULTY
dwigdor@dgp.toronto.edu

Yang, Jishuo (#22)
GRADUATE STUDENT
jishuo.yang@mail.utoronto.ca

Yao, Eric (#25)
UNDERGRADUATE STUDENT
jiaxian.yao@mail.utoronto.ca

Zemel, Richard (#9)
FACULTY
zemel@cs.toronto.edu

Zhao, Jian (#21)
GRADUATE STUDENT
jianzhao@dgp.toronto.edu

www.cs.toronto.edu