

### CONTACT US

For more information on any of these projects:

SARA FRANCA  
Communications & Liaison Officer  
Department of Computer Science  
10 King's College Road, Room 3302  
Toronto, Ontario M5S 3G4  
Telephone: 416-978-3619  
E-mail: sfranca@cs.toronto.edu  
Website: www.cs.toronto.edu

# RESEARCH IN ACTION

RESEARCH HIGHLIGHTS

# CONTENTS

- 1 Basie
- 2 MarkUs Project: Online Marking Made Easy
- 3 Computer-Supported Collaborative Science
- 4 Climate Change Informatics
- 5 The Consequences of Interleaving Input Techniques: Replaying Gestures and Scribbling with Typeset
- 6 NAViGaTOR: Network Analysis, Visualization & Graphing Toronto
- 7 A Bayesian Method for 3D Reconstruction of Macromolecular Structure Using Class Averages from Single Particle Electron Microscopy
- 8 Efficient Software Checkpointing Support for Debugging
- 9 Debugging Domain-Specific Languages
- 10 Dezombify
- 11 Exploring User Fixation With Touch and Mouse Interaction
- 12 RearType: Text Entry Using Keys on the Back of a Device
- 13 Analytic Drawing of 3D Scaffolds
- 14 TimbreMap: a sound-based iPhone application to aid the visually impaired with street map navigation
- 15 Multiscale Symmetric Part Detection and Grouping
- 16 Interactive Performance Control: A New Puppetry of Real and Virtual Robots
- 17 An Unsupervised Model for Text Message Normalization
- 18 Algorithmic Mechanism Design
- 19 Estimating Contact Dynamics from Digital Video
- 20 Side-channel-resilient Cryptography
- 21 Lakeview: Community Education & Landscape Politics
- 22 Runtime Monitoring of Web Service Conversations
- 23 Guided Recovery for Web Service Applications
- 24 Air Computing: "The Cloud" for End Users
- 25 A Matlab Tool-Kit for 3D Cone-Beam Computed Tomography
- 26 DISHA: Multiple Mice-enabled Computer Aided Learning for Children
- 27 Ambulatory Gait Analysis
- 28 Stereo Voice Conferencing SpeakerPhone
- 29 Towards Computational Fluoroscopy
- 30 Using Matrices to Model Symbolic Relationships
- 31 Combining Dynamic and Static Analysis for Program Verification
- 32 Making Physically-Simulated Characters Walk Like People
- 33 Technology in Support of Healthy Aging: Innovations in the Design of Electronic Cognitive Prostheses
- 34 Regret-based Preference Elicitation for Recommender Systems
- 35 Exploiting N-Gram Analysis to Predict Operator Sequences
- 36 Rapidly Adaptive Servers with SnowFlock
- 37 The Model Management Tool Framework
- 38 Adding the Easy Button to the Cloud with SnowFlock and MPI
- 39 Translation from JSCOOP to Multi-threaded Java
- 40 Hyperflow: A Distributed Controller Architecture for OpenFlow Networks
- 41 Spam Detection in IP Telephony
- 42 Lockr: Better Privacy for Social Networks
- 43 Automatic Road Detection in High Resolution Aerial Images
- 44 Gossiping in One-Dimensional Ad Hoc Radio Networks
- 45 Agile Software Process Improvement: Shift of Concerns from Process to Fragment Improvement
- 46 Intelligent Fall Detection and Emergency Response System
- 47 Towards Dynamic End-to-End Resource Allocation in Shared Data Centers
- 48 Making Friends While Choosing Better Books
- 49 Smokebean: Capturing Dragonboat Paddler Work Output for Performance Measurement
- 50 Genetic Variation Discovery
- 51 Relative Approximation to the Singular Values and Applications
- 52 Probabilistically Modeling Energy Consumption in Buildings

## MESSAGE from the chair

The Department of Computer Science (DCS) has been a leading international research department for over forty years. Our roughly fifty research faculty and 300 graduate students are engaged in cutting-edge investigations at the forefront of the dramatic changes in the development and uses of information and communications technology. We also place great emphasis on exposure to, and direct engagement in, research projects among our undergraduates, integrating inquiry-based learning deeply into our curriculum.

We are pleased to highlight a small cross-section of the many fascinating research projects currently underway in DCS. The research described in this booklet gives a sense of the fundamental advances in computing taking place in the department, from network design, virtualization techniques, and cryptography, to new tools for software modeling and verification or interacting with computing devices. It also demonstrates the wide-ranging impact of computer science on areas as diverse as animation; social networks; medical diagnostics; biological discovery; assistive technologies for children, the elderly, and people with physical impairments; and much more. The projects described cover the entire spectrum from applications-oriented research that brings new insights to practical problems, to groundbreaking theoretical research that influences the entire discipline.

Much of our innovative research is shaped by collaborations with research partners in industry and in other academic disciplines – in the sciences, health sciences, social sciences and humanities. Connections and collaborations with industry are especially important to ensure that our research has genuine impact. We're proud of our deep interdisciplinary connections and our strong ties with the industrial and business communities, but are always striving to do more. I encourage potential partners to explore the breadth and depth of the internationally renowned research in DCS, and propose novel ways to benefit from our expertise, creativity and drive.

*Craig Boutilier  
Professor & Chair*



1

## Basie

Basie is a web-based software project management forge that offers multiple projects per forge, pluggable authentication, role-based access control, web-based administration, and a REST API, and provides wikis, repository browsers, mailing lists, ticketing, full-featured search, and a graphical status dashboard. It is 100% open source, self-hosting, and easy to set up and administer.



### PROJECT TEAM

**William Konrad**, Undergraduate Student  
**Ian Lienert**, Undergraduate Student  
**Zuzel Vera Pacheco**, Graduate Student  
**Greg Wilson**, Faculty

2

## MarkUs Project: Online Marking Made Easy

MarkUs is an open-source tool which recreates the ease and flexibility of grading assignments with pen on paper, within a web application. It also allows students and instructors to form groups, and collaborate on assignments. As students submit their work, MarkUs keeps track of the versions they submit. Graders annotate students' code and assign grades directly in the web application. Instructors can monitor the progress of the graders and can easily release the results to the students.



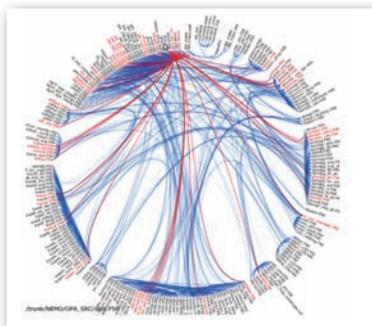
### PROJECT TEAM

**Karen Reid**, Faculty  
**Mike Conley**, Graduate Student  
**Severin Gehwolf**, Undergraduate Student

3

## Computer-Supported Collaborative Science

Our tools support computational scientists as they develop scientific software, document their findings and share their results with others. TracSNAP uses social network analysis to help large teams of scientists to coordinate their work when developing shared code for scientific application; MyeLink supports the use of wikis as open laboratory notebooks by comparing the structure and content of different wiki pages; Breadcrumbs allows scientists to save and share graphs of related web pages when browsing for related work; and VizExpress allows scientists to insert graphs and tables into wiki pages.



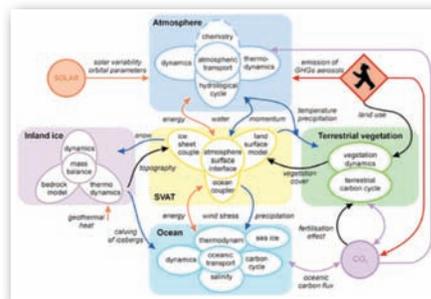
### PROJECT TEAM

**Steve Easterbrook**, Faculty  
**Ainsley Lawson**, Undergraduate Student  
**Sarah Strong**, Undergraduate Student  
**Brent Mombourquette**, Undergraduate Student  
**Samar Sabie**, Undergraduate Student  
**Maria Yancheva**, Undergraduate Student

4

## Climate Change Informatics

Our research group is exploring the role of software in understanding and responding to the threat of climate change. This includes analysis of the software engineering practices of climate scientists as they build and validate complex earth system simulation models; the use of coordination and social networking tools to foster the sharing of code and data across inter-disciplinary scientific teams; the role of software in developing and deploying tools for education and communication about the latest findings in climate science; and support for effective decision-making by individuals, communities and governments as we transition to a carbon-neutral society.



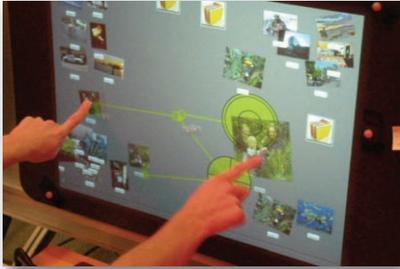
### PROJECT TEAM

**Steve Easterbrook**, Faculty  
**Jon Pipitone**, Graduate Student  
**Alicia Grubb**, Graduate Student  
**Jorge Aranda**, Graduate Student  
**Jennifer Horkoff**, Graduate Student  
**Carolyn MacLeod**, Graduate Student  
**Jonathan Lung**, Graduate Student

5

## The Consequences of Interleaving Input Techniques: Replaying Gestures and Scribbling with Typeset

Interleaving between multiple input techniques is not well explored beyond the time-cost of switching. We present two prototypes that explore this interplay. In normal application use, previous gestural actions may be re-used, much like words are re-used and gain new meaning in conversation. We present Replaying Gesture Macros, work on recording, replaying and modifying multi-touch gestures. Typeset text via keyboard input is usable but confined to rigid, linear textboxes, while written text from stylus entry is slower and harder to read and re-use. We present Scribbling with Typeset, a prototype that combines the best expressive properties of these two modes.



PROJECT TEAM

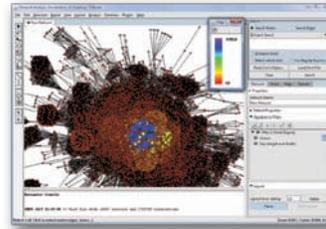
**Dustin Freeman**, Graduate Student  
**Ravin Balakrishnan**, Faculty

6

## NAViGaTOR: Network Analysis, Visualization & Graphing Toronto

NAViGaTOR is a powerful graphing application for the 2D and 3D visualization of biological networks. NAViGaTOR includes a rich suite of visual mark-up tools for manual and automated annotation, fast and scalable layout algorithms, and OpenGL hardware acceleration to facilitate scalable visualization of large graphs. Publication-quality images can be rendered through SVG graphics export, and many community-developed data formats support data import/export (PSI-XML, BioPax, and GML).

NAViGaTOR's ability to handle larger datasets is facilitated through optimized layout algorithms, hardware-based graphics acceleration, and a reduced memory footprint relative to other software. The NAViGaTOR user interface includes unique tools to help simplify the "hairball", which is a common challenge in many PPI networks.



PROJECT TEAM

**Igor Jurisica**, Faculty and and Ontario Cancer Institute  
**Kevin R. Brown**, Ontario Cancer Institute  
**David Otasek**, Ontario Cancer Institute  
**Muhammad Ali**, Ontario Cancer Institute  
**Michael McGuffin**, Ecole de technologie superieure  
**Wing Xie**, Ontario Cancer Institute  
**Baiju Devani**, Laboratory of Integrative Motor Behavior  
**Ian Lawson van Toch**, Cancer Informatics Fund (*Deceased*)

7

## A Bayesian Method for 3D Reconstruction of Macromolecular Structure Using Class Averages from Single Particle Electron Microscopy

The goal of our project is to develop a family of methods capable of constructing three dimensional structure models of target macromolecules (typically large protein complexes) given a set of two-dimensional tomographic images produced by electron microscopy. The input images are somewhat similar in information content to clinical x-ray or CAT-scan images; however, they have extremely low resolution and represent 'views' of the protein from unknown orientations. Our Bayesian method performs three-dimensional model building and generates a structure that best explains the observed electron cryo-microscopy data. Validation on two experimental and one simulated data set shows good reconstruction in all three cases.

PROJECT TEAM

**Navdeep Jaitly**, Graduate Student  
**Marcus Brubaker**, Graduate Student  
**John Rubinstein**, Faculty, Biochemistry and Hospital for Sick Children  
**Ryan Lilien**, Faculty

[www.cs.toronto.edu](http://www.cs.toronto.edu)

8

## Efficient Software Checkpointing Support for Debugging

Checkpointing is a process of taking program snapshots to facilitate program recovery from any unexpected events. Based on specialized hardware, most existing checkpointing approaches create snapshots on persistent storage, causing prohibitive overhead. In contrast, we propose a software-only solution that has no need for specialized checkpointing hardware. It is enabled by compiler analysis that will selectively backup only changed program states into memory. In addition, optimizations will guarantee we only backup the absolute-minimal amount. As a first application, we implement the checkpointing support in a debugger, enabling its roll-back capability to any designated program points without noticeable delays.

PROJECT TEAM

**Greg Steffan**, Faculty, Electrical & Computer Engineering  
**Cristiana Amza**, Faculty, Computer Science and Electrical & Computer Engineering  
**Chengyan (Chuck) Zhao**, Graduate Student

[www.cs.toronto.edu](http://www.cs.toronto.edu)

## Debugging Domain-Specific Languages

Much research focuses on the techniques, tools and benefits of domain-specific language creation. However, since software maintenance is an important part of software projects, we should consider the maintainability of domain-specific language programs. These programs differ in many ways from their general-purpose language counterparts; we should not expect results applicable to one set to generalize to the other. We have done a case study of the maintenance of domain-specific language programs in industry.

```

validates_presence_of :height_id, :hair_color_id, :ethnicity_id
validates_presence_of :first_name, :last_name, :postal_code, :city, :country
validates_presence_of :headline, :description

validates_presence_of :login, :email
validates_presence_of :password, :if => :password_required?
validates_presence_of :password_confirmation, :if => :password_required?
validates_length_of :password, :within => 4..40, :if => :password_required?
validates_confirmation_of :password, :if => :password_required?
validates_length_of :login, :within => 3..40
validates_length_of :email, :within => 3..100
validates_uniqueness_of :login, :email, :case_sensitive => false

before_save :encrypt_password
after_create :create_details

has_one :detail
has_one :avatar

has_many :posts do
  def recent
    find(:all, :order => 'created_at desc', :limit => 1)
  end
end

```

PROJECT TEAM

**Greg Wilson**, Faculty  
**Aran Donohue**, Graduate Student

## Dezombify

Mixed lighting is a problem faced by photographers everywhere. Yellowish incandescent light make people look like they have jaundice. Harsh blue light from camera flashes make people appear pale like zombies. Through chromatic adaptation, human perception allows us to see the true skin tones of people regardless of the light colour illuminating the scene. Virtually all digital cameras lack this corrective process.

Dezombify is a semi-automated process for applying digital chromatic adaptation to large photo collections. It uses face detection to identify skin tones and recolours photographs so people appear less like the undead.



PROJECT TEAM

**Greg Wilson**, Faculty  
**Jason Montojo**, Graduate Student

## Exploring User Fixation With Touch and Mouse Interaction

As touch technology matures, it is becoming commonplace in a variety of computational media including traditional desktop environments where the mouse and keyboard have previously been the primary input mechanisms. We explore how direct and absolute touch input mechanisms can co-exist with indirect and relative mouse/keyboard input. Further, we have conducted observational studies on how users interleave between these two input styles when performing typical day-to-day tasks. We also observed if and how users fixate on a particular input style even if the other is known to be functionally superior. The insights gained from our studies suggest ways in which touch plus mouse/keyboard interfaces of the future can be improved to support multiple input styles used in tandem.

PROJECT TEAM

**Dominika Ruskowski**, Graduate Student  
**Ravin Balakrishnan**, Faculty  
**Jeremy Bernholtz**, Cornell University

## RearType: Text Entry Using Keys on the Back of a Device

RearType is a text input system for larger form-factor mobile devices such as Tablet PCs, which uses normal keyboard keys, but on the reverse side of the device. The aim is to 'free up' the front of the device, maximizing the use of a physical display, minimizing the need for an onscreen keyboard and resulting pixel occlusion, and providing a tactile and multi-finger method for text entry – with the potential for knowledge transfer from QWERTY. Using a prototype implementation which includes software visualization of the keys to assist with learning, we conducted a study to explore the initial learning curve for RearType, and contrast performance with a touchscreen soft-keyboard and a regular keyboard. With one hour's training, RearType entry speeds of 15 WPM were statistically indistinguishable from speeds achieved on the touchscreen keyboard, showing potential for such a novel form of text input for tablet devices. This research project is a Microsoft-UofT collaboration.

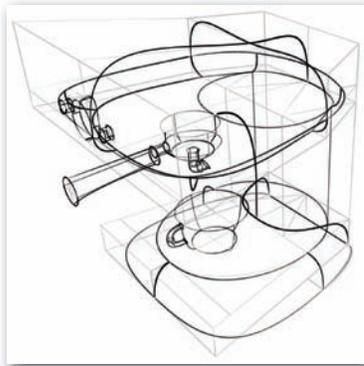
PROJECT TEAM

**James Scott**, Microsoft Research  
**Shahram Izadi**, Microsoft Research  
**Leila Rezai**, Aachen University  
**Dominika Ruskowski**, Graduate Student  
**Xiaojun Bi**, Graduate Student  
**Ravin Balakrishnan**, Faculty

13

## Analytic Drawing of 3D Scaffolds

We have developed a novel approach to inferring 3D curves from perspective drawings in an interactive design tool. Our methods are based on a traditional design drawing style known as analytic drawing, which supports precise image-space construction of a linear 3D scaffold. This scaffold in turn acts as a set of visual constraints for sketching 3D curves.



### PROJECT TEAM

**Karan Singh**, Faculty  
**Ryan Schmidt**, Graduate Student  
**Azam Khan**, Autodesk  
**Gord Kurtenbach**, Autodesk

14

## TimbreMap: a sound-based iPhone application to aid the visually impaired with street map navigation

TimbreMap is an iPhone application that conveys street map patterns to the user using directional, audio feedback from touch interaction, making it a great tool to assist the visually impaired with street map navigation and learning.

### PROJECT TEAM

**Khai Truong**, Faculty  
**Eyal de Lara**, Faculty  
**Alyssa Rosenzweig**, Graduate Student  
**Jing Su**, Graduate Student

15

## Multiscale Symmetric Part Detection and Grouping

Skeletonization algorithms typically decompose an object's silhouette into a set of symmetry-based parts, offering a powerful representation for shape categorization, but require correct figure-ground segmentation. We present a novel approach to recovering and grouping the symmetric parts of an object from a cluttered scene without assuming prior figure/ground segmentation. We start by clustering superpixels at multiple scales into symmetric parts and proceed by grouping the symmetric parts into objects. The resulting framework yields a skeletal approximation that's free of many of the instabilities plaguing traditional skeletons and doesn't require a closed contour, enabling the application of skeleton-based categorization systems to more realistic imagery.

### PROJECT TEAM

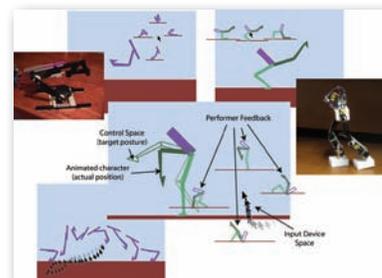
**Alex Levinshtein**, Graduate Student  
**Cristian Sminchisescu**, Bonn University (Germany)  
**Sven Dickinson**, Faculty

16

## Interactive Performance Control: A New Puppetry of Real and Virtual Robots

From video games to feature films, the laws of physics play an increasing role in creating the compelling animation modern audiences demand. Animating agile, physically realistic characters, however, is as much art as it is science, whether the characters are inhabitants of fantastical virtual worlds, real robotic competitors running in a soccer stadium, or mechanical family pets.

Our approach puts the reins directly into the hands of the animator, who becomes an interactive performer, vicariously flexing a character's muscles or motors, moment-to-moment, to make it run and jump about under its own power, a "stringless puppet" in a potential new performance art form. We present a variety of simple but compelling examples of animated characters "performed" in this fashion, and argue the approach to be a viable form of physics-based animation, engaging and entertaining to performer and audience alike.



### PROJECT TEAM

**Joe Laszlo**, Graduate Student  
**Eugene Fiume**, Faculty

## An Unsupervised Model for Text Message Normalization

Cell phone text messages contain an abundance of creative shortenings, for example “How are you tonight?” could be written as “Hw r u tonite”. In order to apply natural language processing techniques for tasks such as machine translation to text messages, the text must first be normalized, that is, given a text message such as “Hw r u tonite”, the standard form “How are you tonight” must be inferred. Drawing on observations about common ways in which words are modified in text messaging, we propose a statistical method for text message normalization. Our method gives performance on par with the best previously reported results on the dataset we use, while also having the potential to be easily applied to new text genres which contain similar abbreviations, such as Twitter.

PROJECT TEAM

**Paul Cook**, Graduate Student  
**Suzanne Stevenson**, Faculty

## Algorithmic Mechanism Design

We consider combinatorial auctions in which objects are sold to selfish bidders, each having a private valuation function that expresses values for sets of objects. In particular, we consider the social welfare objective where the goal of the auctioneer is to allocate the objects so as to maximize the sum of the values of the allocated sets. Our specific interest is to understand the power and limitations of “simple” allocation algorithms in this regard. We show that monotone greedy algorithms can be made into mechanisms whose price of anarchy (i.e. at every ex-post Nash equilibria) is almost as good as the approximation that the greedy allocation achieves (disregarding selfish interest). In contrast to this positive result on the price of anarchy, we also consider the limitations of greedy truthful mechanisms, where we model the notion of greediness with a broad class of algorithms known as priority algorithms. For single minded bidders, there is a known truthful deterministic mechanism (Lehmann, O’Callaghan and Shoham) using a greedy allocation that achieves an  $O(\sqrt{m})$  approximation ratio to the social welfare. We show that the approximation ratio for any truthful priority algorithm must be asymptotically much worse than what is achievable without the truthfulness requirement.

PROJECT TEAM

**Allan Borodin**, Faculty  
**Brendan Lucier**, Graduate Student

## Estimating Contact Dynamics from Digital Video

We want to infer 3D human motion and peoples’ interactions with the world from images. Our recent work has shown how one can infer the internal and external forces acting on an articulated body. In particular, we present a formulation whereby one can estimate the internal forces (muscle actuations) and external forces (gravity and contact) acting on an articulated body, detect contact between the body and external surfaces, and infer the parameters of a contact model (e.g., surface geometry and spring parameters). This method will help learn physics-based models for computer animation, and provide the basis for visual analysis of human motion and scene interactions.



PROJECT TEAM

**Marcus Brubaker**, Graduate Student  
**Leon Sigal**, Postdoctoral Fellow  
**David Fleet**, Faculty

## Side-channel-resilient Cryptography

In cryptography, it is typically assumed that software runs on devices that are secure. That is, the devices on which cryptographic software runs are modeled as “black boxes”. In reality, the security of many encryption systems has been broken by making simple physical observations, known as side-channel attacks, on a device’s environment. In side-channel attacks, an adversary obtains information about the internal state of a device by measuring such things as power consumption or emitted radiation. In this work, we show how to compute every keyed primitive (such as pseudorandom functions or signature schemes) in a way that is resilient to such attacks.

PROJECT TEAM

**Charles Rackoff**, Faculty  
**Ali Juma**, Graduate Student  
**Yevgeniy Vahlis**, Graduate Student

## Lakeview: Community Education & Landscape Politics

UofT's Centre for Landscape Research (CLR) explores new directions for academic practice by enabling landscape architectural thinking across a whole community. CLR uses innovative design software that allow the Lakeview Ratepayers to examine, question, research, and creatively vision possible landscape futures and convince politicians to act on community design ideas.

PROJECT TEAM  
**John Danahy**, Faculty

## Runtime Monitoring of Web Service Conversations

Web services are distributed processes that dynamically discover properties of other web services. We propose to use runtime monitoring of conversations between partners as a means of checking behavioural correctness of the entire web service system. Specifically, as part of an on-going collaboration with IBM Toronto Lab, we have identified a subset of UML 2.0 Sequence Diagrams as a property specification language and showed that it can capture interesting safety ("the system will not engage in a certain negative behaviour") properties. By transforming these diagrams into finite-state automata, we give semantics to our specification language and enable conformance checking of finite execution traces against the specification. We also show how our language can be used to specify property templates.

### PROJECT TEAM

**Jocelyn Simmonds**, Graduate Student  
**Marsha Chechik**, Faculty  
**Shiva Nejati**, Graduate Student  
**Yuan Gan**, Graduate Student  
**Bill O'Farrell**, IBM  
**Elena Litani**, IBM  
**Julie Waterhouse**, IBM

## Guided Recovery for Web Service Applications

Recovery frameworks offer mechanisms that can be used to help applications recover from errors as they are detected. Current recovery frameworks for web services have various limitations: 1) most frameworks assume that errors only appear on interactions with partner services, and 2) recovery strategies usually need to be manually specified. We propose a more versatile framework that will generate recovery sequences, with the goal of returning the system to a stable state. If no such sequences are available, execution terminates. During the course of this project, we have identified various heuristics that allow the automated identification of candidate stable states. Recovery plans are generated using an AI planner. Our poster will present some initial results.

PROJECT TEAM  
**Jocelyn Simmonds**, Graduate Student  
**Marsha Chechik**, Faculty  
**Shoham Ben-David**, Postdoctoral Fellow  
**Bill O'Farrell**, IBM  
**Elena Litani**, IBM  
**Leho Nigul**, IBM

## Air Computing: "The Cloud" for End Users

While cloud computing has been widely adopted by enterprises, it has had a far less impact on the consumer market. This is due to three main reasons: 1) cloud computing is not cost effective for end users, 2) it is not suited to highly interactive desktop applications, 3) ownership and privacy concerns. Therefore, we propose Air Computing, a new paradigm that seamlessly combines users' own devices "in the air" with server resources in the cloud. This provides average users with the flexibility and simplicity to create a computing infrastructure suited to her own needs – that is, the user can easily build a "private" cloud using her own devices, and optionally integrate with cloud server resources to augment her personal cloud.

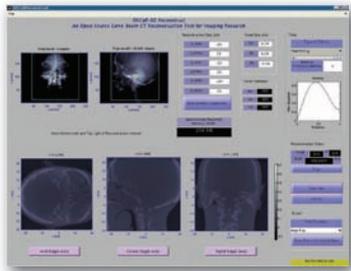


PROJECT TEAM  
**Weihan Wang**, Graduate Student, Electrical & Computer Engineering  
**Cristiana Amza**, Faculty, Computer Science and Electrical & Computer Engineering

## A Matlab Tool-Kit for 3D Cone-Beam Computed Tomography

Advances in x-ray detectors within the last ten years has improved the quality of computerised tomography (CT) volumetric imaging systems significantly. In particular, data acquisition using high resolution flat-panel x-ray detectors permits cone-beam computerised tomography (CBCT). Earlier technologies relied on acquiring a large number of scans of horizontal slices using parallel-or fan-beam geometries and performing two-dimensional reconstructions over each slice. CBCT systems capitalise on rapid acquisition of circular cone-beam data in a single scan and subsequently use the Feldkamp-Davis-Kress (FDK) formulation for three-dimensional reconstructions. We present some of our Matlab implementation that bundles FDK-type schemes with

the intention of providing an open-source research tool. Our software is called OSCaR: An Open-Source Cone-Beam CT Reconstruction Tool for Imaging Research.



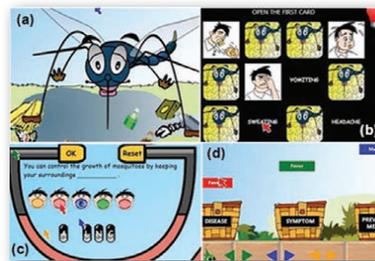
### PROJECT TEAM

**Nargol Rezvani**, Graduate Student  
**Dhavid Aruliah**, University of Ontario Institute of Technology  
**Ken Jackson**, Faculty  
**Douglas Moseley**, Princess Margaret Hospital  
**Jeffrey Siewerdsen**, Johns Hopkins University

## DISHA: Multiple Mice-enabled Computer Aided Learning for Children

Global health statistics show a staggering amount of deaths and disabilities due to preventable diseases including respiratory infections, malaria, diarrhea and tuberculosis that almost exclusively impact the developing world. DISHA (DISEase and Health Awareness) is a collaborative platform for public health information regarding the symptoms, prevention and cure of the disease malaria, for children in low-income regions using multiple mice. DISHA uses a narrative-interactive loop format of story-telling followed by Q&A. The Q&A enforces collaboration, competition and turn taking as machine-induced interactive modules, following a game based approach. Standard design principles of gaming including goal-based progression, positive reinforcement, Same

Language Subtitling (SLS) and color and shape-coded mouse pointers were employed.



### PROJECT TEAM

**Mohit Jain**, Graduate Student  
**Aakar Gupta**, Intern, Microsoft Research India  
**Navkar Samdaria**, Intern, Microsoft Research India  
**Praveen Shekhar**, Intern, Microsoft Research India  
**Joyojeet Pal**, Postdoctoral Researcher, University of Washington

## Ambulatory Gait Analysis

Working with researchers at Keio University in Japan we have developed a small 3D acceleration sensing package that can be embedded in clothing and shoes. Using this package, belt, slipper, and shoe-based ambulatory (measurement while walking) gait analysis prototypes will be demonstrated along with visualizations of gait collected by the prototypes. The ambulatory gait analysis is intended for a wide range of purposes, including assessment of falls risk in nursing home residents.

### PROJECT TEAM

**Mark Chignell**, Faculty, Computer Science and Mechanical & Industrial Engineering  
**Anita Ko**, Graduate Student, Mechanical and Industrial Engineering

## Stereo Voice Conferencing SpeakerPhone

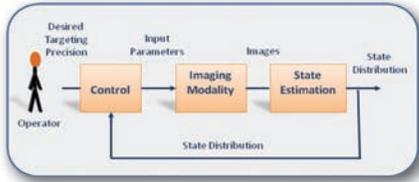
This demo will highlight speakerphone and stereo echo cancellation features that have recently been added to Vocal Village spatialized audio conferencing. Two models of speakerphone will be demonstrated based on a Nortech computer and a Beagleboard respectively.

### PROJECT TEAM

**Mark Chignell**, Faculty, Computer Science and Mechanical & Industrial Engineering  
**Hao Shi**, Graduate Student, Mechanical and Industrial Engineering

## Towards Computational Fluoroscopy

Several complex image-guided procedures, particularly in interventional cardiology and neuroradiology, and in radiotherapy, involve high dose rates and long treatment times. There is growing concern regarding the dose delivered during these procedures as the ionizing nature of X-rays leads to immediate and long term damaging effects. This work introduces a framework for balancing tradeoffs between targeting precision and dose by dynamically adapting imaging parameters using feedback of geometric performance. In this framework, the operator specifies a desired targeting precision for a procedure and the system adapts itself to maintain this level of targeting precision using minimal dose.

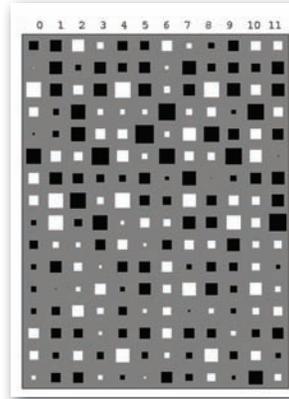


PROJECT TEAM

**Eugene Fiume**, Faculty  
**Sami Siddique**, Graduate Student  
**David Jaffray**, PMH

## Using Matrices to Model Symbolic Relationships

We describe a way of learning matrix representations of objects and relationships. The goal of learning is to allow multiplication of matrices to represent symbolic relationships between objects and symbolic relationships between relationships, which is the main novelty of the method. We demonstrate that this leads to excellent generalization in two different domains: modular arithmetic and family relationships. We show that the same system can learn first-order propositions such as (2, 5) in +3 or (Christopher, Penelope) in “has wife”, and higher-order propositions such as (3, +3) in “plus” and (+3, -3) in “inverse” or (“has husband”, “has wife”) in “higher oppsex”. We further demonstrate that the system understands how higher-order propositions are related to first-order ones by showing that it can correctly answer questions about first-order propositions involving the relations “+3” or “has wife” even though it has not been trained on any first-order examples involving these relations.



PROJECT TEAM

**Ilya Sutskever**, Graduate Student  
**Geoff Hinton**, Faculty

## Combining Dynamic and Static Analysis for Program Verification

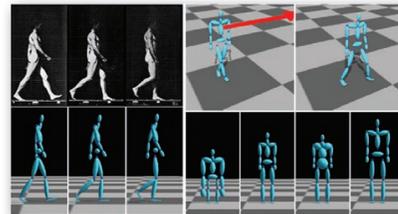
We contribute an algorithm for the verification of concurrent programs. Our algorithm combines the power of concrete execution - dynamic analysis - to find errors with the power of abstraction - static analysis - to prove the unreachability of errors. This amalgamation allows results from dynamic analysis to guide the static analysis and vice versa. The approach is sound, but incomplete due to the undecidability of the problem in general.

PROJECT TEAM

**Aws Albarghouthi**, Graduate Student  
**Marsha Chechik**, Faculty  
**Ou Wei**, Alumnus  
**Arie Gurfinkel**, Carnegie Mellon University

## Making Physically-Simulated Characters Walk Like People

Creating character animation through interactive physical simulation offers great potential for increased realism and unpredictability in interactive games. However, control algorithms which determine virtual muscle activations depending on the environment are extremely difficult to define, even for basic actions such as walking. Drawing from biomechanical knowledge, we show an off-line optimization process which can be used to create walking control algorithms. Our results capture a number of properties of natural human walking, and are robust to minor environmental disturbances.



PROJECT TEAM

**Jack Wang**, Graduate Student  
**David Fleet**, Faculty  
**Aaron Hertzmann**, Faculty

## Technology in Support of Healthy Aging: Innovations in the Design of Electronic Cognitive Prostheses

The Technologies for Aging Gracefully lab (TAG lab) carries out the human-centred design of assistive, rehabilitative, and health-preserving technology for individuals throughout the life course, with an emphasis on senior citizens and people suffering from cognitive and communication impairments. Three research projects will be featured:

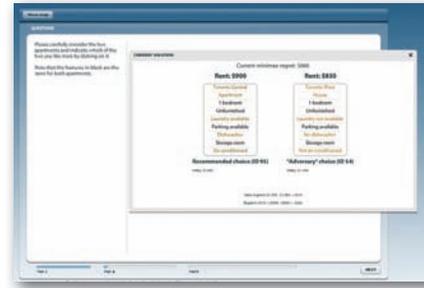
1. Lifelogging and digital storytelling to enhance recall and reminiscence in Alzheimer's disease.
2. Context-aware mobile devices to facilitate word and name recall and speech for anomic aphasia.
3. Cognitively-and socially-stimulating gaming environments to enhance cognitive reserve in seniors.

### PROJECT TEAM

**Ron Baecker**, Faculty  
**Mike Massimi**, Graduate Student  
**Mike Wu**, Graduate Student  
**Velian Pandeliev**, Graduate Student  
**Miller Peterson**, Graduate Student  
**Nick Shim**, Graduate Student  
**Kevin Tonon**, Graduate Student  
**Alexander Levy**, Undergraduate Student  
**Aakash Sahney**, Undergraduate Student  
**Garry Ing**, Research Assistant  
**Masashi Crete-Nishihata**, Research Assistant  
**Deb Ptak**, Research Assistant  
**Mark Watson**, Research Assistant

## Regret-based Preference Elicitation for Recommender Systems

Recommender systems can help users navigate product spaces and make decisions involving large sets of alternatives. The nature of web interactions means that recommendations have to be made with very limited information about the user's preferences. In this work we adopt minimax-regret decision criterion for product recommendation under utility uncertainty, where user feedback is incorporated as constraints to the utility function. We show how to produce recommendations in face of feature uncertainty (with user-defined attributes) and how to optimize a set of joint recommendations.

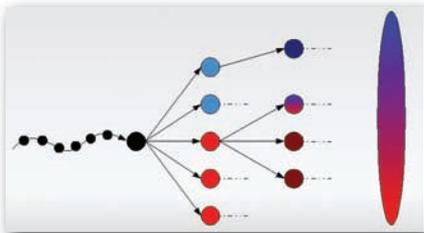


### PROJECT TEAM

**Craig Boutilier**, Faculty  
**Paolo Viappiani**, Postdoctoral Fellow  
**Darius Braziunas**, Graduate Student  
**Kevin Regan**, Graduate Student

## Exploiting N-Gram Analysis to Predict Operator Sequences

N-gram analysis provides a means of probabilistically predicting the next item in a sequence. Due originally to Shannon, it has proven an effective technique for word prediction in natural language processing and for gene sequence analysis. In this project, we investigate the utility of n-gram analysis in predicting operator sequences in plans. Given a set of sample plans, we perform n-gram analysis to predict the likelihood of subsequent operators, relative to a partial plan. We identify several ways in which this information might be integrated into a planner. In this project, we investigate one of these directions in further detail. Preliminary results demonstrate the promise of n-gram analysis as a tool for improving planning performance.



### PROJECT TEAM

**Christian Muise**, Graduate Student  
**Sheila McIlraith**, Faculty  
**Jorge Baier**, Graduate Student  
**Michael Reimer**, Graduate Student

## Rapidly Adaptive Servers with SnowFlock

When users download or store a file, they just use more of the shared pool of bandwidth and disk space. It's a different story for servers that need more CPU or RAM, even with cloud computing. Since booting a new physical or virtual machine (VM) takes minutes, it's hard to maintain the quality of service when a spike in load exceeds available power. Our rapidly adaptive server uses Snowflock's fast VM cloning to transparently spread to more hosts machines as needed, and automatically retreats as possible for lower costs. With new VMs in less than a second, we don't need to anticipate future load or keep pre-booted VMs on standby. Our research investigates the efficiency and performance of our prototype running Apache web server under a heavy and variable load.



### PROJECT TEAM

**Michael Brudno**, Faculty  
**Eyal de Lara**, Faculty  
**Roy Bryant**, Graduate Student  
**Olga Irzak**, Graduate Student

## The Model Management Tool Framework

The Model Management Tool Framework (MMTF) is an Eclipse-based tool framework for software model management. At a glance, MMTF allows users to create tools based on a kind of models (“macromodels”) that can depict a high level view of other models and their relations. Among other applications, this enables users to capture aspects such as partial modeling, deferred design decisions and variability, to check for consistency and validity, and in general to use the macromodels to drive development and capture model evolution. Our current work focuses on using MMTF’s capabilities to capture the development of automotive software in a model-driven production environment.

### PROJECT TEAM

**Marsha Chechik**, Faculty  
**Michalis Famelis**, Graduate Student  
**Rick Salay**, Graduate Student  
**Qiyu Zhu**, Undergraduate Student  
**Robert Baillargeon**, Panasonic Corporation

## Adding the Easy Button to the Cloud with SnowFlock and MPI

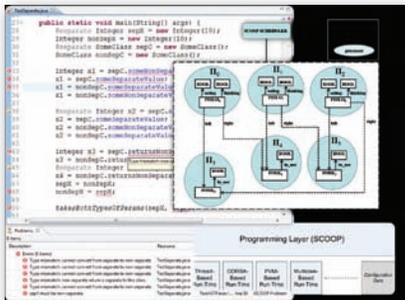
Cloud computing promises to provide researchers with the ability to perform parallel computations using large pools of virtual machines (VMs), without facing the burden of owning or maintaining physical infrastructure. However, with ease of access to hundreds of VMs, comes also an increased management burden. Cloud users today must manually instantiate, configure and maintain the virtual hosts in their cluster. They must learn new cloud APIs that are not germane to the problem of parallel processing. And in the world of high performance computing (HPC) they must manage their applications across all these VMs. We propose the use of SnowFlock’s rapid VM cloning ability with the widely-used HPC standard of MPI.

### PROJECT TEAM

**Michael Brudno**, Faculty  
**Eyal de Lara**, Faculty  
**Philip Patchin**, Graduate Student  
**Andres Lagar-Cavilla**, Alumnus

## Translation from JSCOOP to Multi-threaded Java

SCOOP is a minimal extension to the sequential object-oriented programming model for concurrency. The extension consists of one keyword (separate) that avoids explicit thread declarations, synchronized blocks, explicit waits, and eliminates data races and atomicity violations by construction through a set of compiler rules. It attempts to guarantee fairness via use of a global scheduler. We are developing a new implementation of SCOOP for Java, called JSCOOP. JSCOOP introduces a new set of annotations modeled after SCOOP keywords, as well as several core library classes which provide the support necessary to implement the SCOOP semantics. A prototype Eclipse plug-in allows for the creation of JSCOOP projects. The fact that JSCOOP programs are annotated Java programs means we can exploit existing Java formal methods tools, particularly JML and Java Pathfinder.

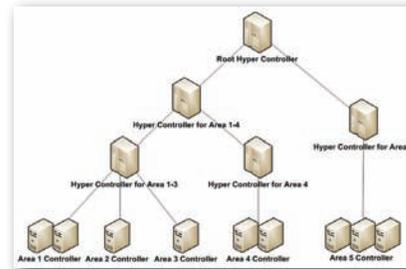


### PROJECT TEAM

**Faraz Torshizi**, Graduate Student  
**Marsha Chechik**, Faculty  
**Cameron Gorrie**, Undergraduate Student

## Hyperflow: A Distributed Controller Architecture for OpenFlow Networks

Today, it is very hard to innovate in networks. OpenFlow is an open standard which enables innovation by moving the control functionality out of the switches to a centralized controller which programs the switches to forward packets. Currently, a major limitation of OpenFlow is its simplifying assumption of a single central controller. This is an obstacle to its deployment in large-scale networks, due to the robustness, availability, and scalability requirements. In this work, we present HyperFlow, a distributed control plane architecture for OpenFlow which does not require any changes to the OpenFlow standard. It provides network



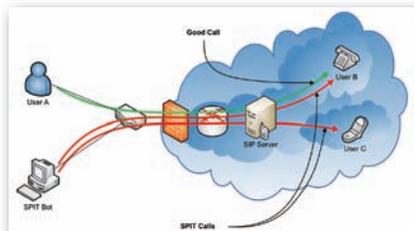
operators with a high flexibility to tune the control plane according to their performance and availability requirements.

### PROJECT TEAM

**Yashar Ganjali**, Faculty  
**Amin Tootoonchian**, Graduate Student

## Spam Detection in IP Telephony

In this project we address the problem of detecting and preventing Spam over Internet Telephony (SPIT). As price of Voice Over IP (VoIP) is going down, SPIT is expected to become a serious problem in the near future. It has the potential to become even a bigger problem than email spam, because recipients will be disturbed by each received SPIT call. We are exploring properties of SPIT calls from conversations to detect them and use these properties to prevent new SPIT calls before interrupting users. Also we are trying to collect a data set containing phone calls information such as caller and callee identification in order to validate our results.



### PROJECT TEAM

**Yashar Ganjali**, Faculty  
**Hossein Kaffash Bokharai**, Graduate Student  
**Alireza Sahraei**, Graduate Student

## Lockr: Better Privacy for Social Networks

One of the biggest challenges seen with the explosion of Online Social Networking (OSN) sites is ensuring the privacy of users' data. OSN sites require users to sign up and define a network of friends, and currently, they provide no or minimum privacy guarantees. We have developed a system, which we call Lockr, to enhance the privacy of personal information shared on social sites. Lockr separates the social networking information from content delivery. The system allows individual users to manage their own social networks and filter outside access to their private data, based on social relationship status.

While our original Lockr implementation made use of a Facebook application to manage a user's social network, our latest version of Lockr is simply a Firefox extension that does not rely on any third-party site. It supports a user's social network



locally and on the user's server of choice. As a proof of concept, we have integrated Lockr with Flickr to show its simplicity and to demonstrate its benefits.

### PROJECT TEAM

**Yashar Ganjali**, Faculty  
**Linda Enciu**, Undergraduate Student  
**Stefan Saroiu**, Microsoft  
**Alec Wolman**, Microsoft  
**Amin Tootoonchian**, Graduate Student

## Automatic Road Detection in High Resolution Aerial Images

At present, the process of creating and updating road maps is costly and time consuming because it is largely performed by hand. The aim of this project is to build a system that automatically detects roads in high resolution aerial images. Machine learning approaches are particularly well suited to this task because expert labeled data in the form of aligned road maps and satellite images is abundant. We scale up standard machine learning techniques to massive amounts of data by performing our computations on graphics processors. The preliminary detection results on a challenging urban dataset are very promising.

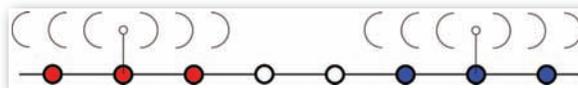


### PROJECT TEAM

**Volodymyr Mnih**, Graduate Student  
**Geoff Hinton**, Faculty

## Gossiping in One-Dimensional Ad Hoc Radio Networks

Consider a wireless radio network consisting of processors located along a line. Examples of networks that can be modelled in this way include electronic surveillance located along a border fence and cars travelling along a highway. In the gossiping task, each processor in the network has a message that must be sent to all of the others. Gossiping is a fundamental communication task in networks and is often used as a building block to complete more complex tasks. We present algorithms for the gossiping task in static and mobile ad hoc radio networks. We also discuss some theoretical limitations of various algorithm designs.



### PROJECT TEAM

**Avery Miller**, Graduate Student  
**Faith Ellen**, Faculty

## Agile Software Process Improvement: Shift of Concerns from Process to Fragment Improvement

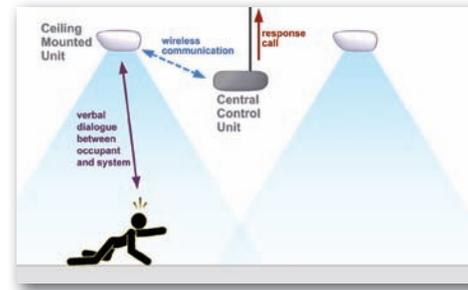
Despite the importance of Software Process Improvement (SPI), many software companies are not willing to get involved in such initiatives, as they are commonly perceived to be long-term, costly, and risky. Agile SPI frameworks have been proposed to overcome these problems. In this research we focus on goal-oriented modeling of software processes for application in agile SPI. Unlike conventional SPI modeling approaches that try to model the whole process of software development, our approach is fragment-based and focuses on modeling and analyzing the motivational forces that lead to the success or failure of specific software development practices. We believe that limiting the scope of SPI from process to fragment improvements can decrease the cost of SPI initiatives, and increase their practical viability for many software companies.

### PROJECT TEAM

**Hesam Chiniforooshan**, Graduate Student  
**Eric Yu**, Faculty, Computer Science and the Faculty of Information

## Intelligent Fall Detection and Emergency Response System

In an attempt to improve safety in the home, there have been several attempts to create a system that allows older adults to call for help in emergency situations. The research team at IATSL has developed an automated emergency response system. The system operates through one or more Ceiling Mounted Units and a Central Control Unit. One Central Control Unit is installed in each home and is responsible for coordinating the Ceiling Mounted Units and relaying communications to the outside world. The use of artificial intelligence allows the system to work entirely autonomously, requiring no input from the user.



### PROJECT TEAM

**Jasper Snoek**, Graduate Student  
**Babak Taati**, Intelligent Assistive Technology and Systems Lab  
**Jen Boger**, Intelligent Assistive Technology and Systems Lab  
**Alex Mihailidis**, Faculty, Computer Science and Occupational Therapy

## Towards Dynamic End-to-End Resource Allocation in Shared Data Centers

With the emerging trend towards server consolidation in large data centers, techniques for dynamic resource allocation for performance isolation between applications become increasingly important. We study the end-to-end resource allocation problem, which involves CPU scheduling, database buffer pool partitioning, and the storage bandwidth partitioning among applications. Our gray-box modeling approach can effectively find a near-optimal configuration which can guarantee per application Service Level Objectives (SLOs) as well as achieve good resource usage.

### PROJECT TEAM

**Jin Chen**, Graduate Student  
**Gokul Soundararajan**, Graduate Student,  
 Electrical & Computer Engineering  
**Cristiana Amza**, Faculty, Computer Science and  
 Electrical & Computer Engineering

## Making Friends While Choosing Better Books

Imagine a software that understands your preferences and, based on them, recommends potential things you would enjoy. In collaboration with the Indigo bookstore we worked on a system to help people pick better books as well as help them connect to other book lovers who share similar interests. Accordingly, we built two systems. One that recommends books to users of Indigo's online store: the system makes use of past user preferences to recommend new books. For the second project, we match people based on their previous book interests using a novel technique.

### PROJECT TEAM

**Arnold Binas**, Graduate Student  
**Laurent Charlin**, Graduate Student  
**Alex Levinshstein**, Graduate Student  
**Maksims Volkovs**, Graduate Student  
**Greg Wilson**, Faculty

## Smokebean: Capturing Dragonboat Paddler Work Output for Performance Measurement

We have developed a device to measure energy output of an athlete in a variety of paddling sports, with dragonboat as our current focus. The device fills a current void in the solution space; namely, performance measurement tools for paddling sports that can be used in situ, and not in a simulated or training pool environment. The device incorporates a number of force and acceleration sensors to determine the work done by the paddler, as well as tracking the motion to enable athlete technique analysis. The application opens up a variety of new methods and opportunities for athlete performance measurement, which is a critical component of any training regime.

### PROJECT TEAM

**Paul Haist**, Undergraduate Student in Electrical & Computer Engineering  
**Phil Lam**, Graduate Student in Mechanical & Industrial Engineering  
**Mark Chignell**, Faculty, Computer Science and Mechanical & Industrial Engineering

## Genetic Variation Discovery

Genetic variation makes individuals unique not only in their appearance, but also in vital health determinants like susceptibility to disease. We will be showcasing a number of computational tools devised to uncover genetic variation using data from the newest generation of sequencing machines.

### PROJECT TEAM

**Michael Brudno**, Faculty  
**Marc Fiume**, Graduate Student  
**Paul Medvedev**, Graduate Student  
**Misko Dzamba**, Undergraduate Student  
**Adrian Dalca**, Graduate Student  
**Seunghak Lee**, Carnegie Mellon University

## Relative Approximation to the Singular Values and Applications

Advances in data collection and storage capabilities during the past decades have led to an information overload in most sciences. Researchers working in domains as diverse as engineering, biology, economics, and information systems, face larger and larger observations and simulations on a daily basis. Such data sets, in contrast with smaller, more traditional data sets that have been studied extensively in the past, present new challenges in data analysis. Traditional statistical methods break down partly because of the increase in the number of observations. A powerful tool for reducing the dimensionality of the data is the Singular Value Decomposition (SVD). In this project, we develop approximation algorithms for the SVD using the method of random projections.

### PROJECT TEAM

**Avner Magen**, Faculty  
**Anastasios Zouzias**, Graduate Student

## Probabilistically Modeling Energy Consumption in Buildings

Buildings account for a large percentage of energy consumption in the developed world, so they are an important focus for any large-scale plan to cut energy use. Unfortunately, each building has a unique combination of design, location, construction defects, and operational characteristics that make it difficult to predict how much savings a change will produce. We develop a data-driven approach to this problem, combining knowledge about physics of the world with collected energy use and operational data from several buildings. We build a model of energy consumption that is simultaneously physically plausible and that matches real world collected data.

### PROJECT TEAM

**Danny Tarlow**, Graduate Student  
**Andrew Peterman**, Walt Disney Imagineering R&D  
**Benedict R. Schwegler**, Walt Disney Imagineering R&D and Stanford University  
**Christopher Trigg**, Stanford University