



**RESEARCH**

**IN ACTION**



Computer Science  
UNIVERSITY OF TORONTO

## ARTIFICIAL INTELLIGENCE (AI)

### Applying a Naive Bayes Similarity Measure to Word Sense Disambiguation | BOOTH 49 (AI)


We replace the overlap mechanism of the Lesk algorithm with a simple, general-purpose Naive Bayes model that measures many-to-many association between two sets of random variables. Even with simple probability estimates such as maximum likelihood, the model gains significant improvement over the Lesk algorithm on word sense disambiguation tasks. With additional lexical knowledge from WordNet, performance is further improved to surpass the state-of-the-art results.

**Tong Wang**  
tong@cs.toronto.edu

**Prof. Graeme Hirst**  
gh@cs.toronto.edu

---

### Computing High Quality Solutions to Probabilistic Planning Problems | BOOTH 50 (AI)



Planning in Artificial Intelligence is the problem of finding action strategies, typically executed by intelligent software agents, autonomous robots, or unmanned vehicles. Given a description of the world, a set of actions that an agent is capable of performing, and some desired goal or objective, the solution to a planning problem is a policy – a mapping from states to actions – that the agent can execute to achieve the goal. There is often uncertainty in the task either because the world is not modeled precisely, or because action

outcomes are intrinsically uncertain or outside the control of the agent. For example, a slippery road or strong wind can affect the movement of a vehicle. In this work we address the class of Probabilistic Planning problems, where the outcomes of the actions are non-deterministic and follow a probabilistic transition model. In particular, we focus on finding policies that maximize the probability of reaching a prescribed goal. Our algorithm, Prob-PRP, outperforms the state of the art, computing substantially more robust policies orders of magnitude faster than the state of the art.

**Alberto Camacho**  
acamacho@cs.toronto.edu

**Prof. Sheila McIlraith**  
sheila@cs.toronto.edu

**Christian Muise**  
University of Melbourne

---

### Financial Sentiment Analysis | BOOTH 53 (AI)

Analysts believe ...  
skyrocket ... in other  
news ... Expectations  
were ... Downgrade ...  
Hostile treatment ...  
or ... Earning surprise

There are numerous studies suggesting that published news stories have an important effect on the direction of the stock market, its volatility, the volume of trades, and the value of individual stocks mentioned in the news. There is even some published research suggesting that automated sentiment analysis of news documents, quarterly reports, blogs and/or Twitter data can be productively used as part of a trading strategy. This paper presents just such a family of trading strategies that produced annualized returns of 70.1%, and then uses this application to re-examine some of the tacit assumptions behind how sentiment analyzers


are generally evaluated, in spite of the contexts of their application. This discrepancy comes at a cost.

**Siavash Kazemian**  
kazemian@cs.toronto.edu

**Prof. Gerald Penn**  
gpenn@cs.toronto.edu

**Shunan Zhao**  
szhao@cs.toronto.edu

## Latent Factor Models of Travel Data for Travel Prediction and Analysis | BOOTH 3 (AI)



In this project, we build statistical models of human travel. First, we use those models to predict human travel: we would like to know how likely a person is to go from point A to point B within a certain time interval. Second, we analyze the parameters of our models in order to gain insights into how people travel.

We decompose the likelihood of traveling from A to B into three factors: the desirability of B as a destination, the affinity between source A and destination B, and the individual-varying propensity to travel the distance between A and B. By analyzing the models that we learn from two large datasets, geotagged Flickr photos and tracks of Shanghai taxis, we estimate the desirabilities of destinations on the map and affinities between locations (the desirability of the location is distinct from its popularity, since popularity is in part due to proximity to large population centres), as well as discover clusters of individuals with varying propensities to travel large distances.


We analyze the affinities between locations on the map. We confirm that affinities are higher within national borders than across national borders. More interestingly, we find that affinities tend to be higher than average within some linguistic areas but not others. See [www.cs.toronto.edu/~guerzhoy/humantravel](http://www.cs.toronto.edu/~guerzhoy/humantravel) for details.

**Michael Guerzhoy**, Lecturer  
guerzhoy@cs.toronto.edu

**Prof. Aaron Hertzmann**, Adobe Research  
hertzman@dgp.toronto.edu

---

## Multi-cue Mid-level Grouping | BOOTH 34 (AI)



Region proposals are used in state-of-the-art recognition systems to provide plausible object locations in images of cluttered scenes. Compared with sliding windows, region proposals are richer and dramatically fewer, but they still number in the thousands for images of a few objects. In this project, we leverage mid-level visual knowledge to resolve bottom-up ambiguity at an earlier stage with the aim of producing fewer but higher quality

proposals. We generate regions using a combination of grouping cues like appearance, closure, and symmetry, in which the strengths of one compensate for the weaknesses of others. We demonstrate that our model is mid-level by training it on a limited number of objects and applying it on different objects. In our quantitative evaluation, we measure a positive contribution from each individual cue, and demonstrate improvement on two leading region proposal methods under a limited budget of proposals.

**Tom Lee**  
tshlee@cs.toronto.edu

**Prof. Sven Dickinson**  
sven@cs.toronto.edu

**Prof. Sanja Fidler**  
fidler@cs.toronto.edu

### Neural Network for Professional Basketball | BOOTH 4 (AI)


This is a collaboration with our local NBA team, the Raptors. They are interested in the question of shot prediction. Given all the information about a shot (e.g. who is shooting, where the shot happened), can we accurately predict its outcome. In the past, they have solved this problem by using simple spreadsheets. However, when the factors that we might be interested in increases as we have more ways to extract features for the game, the number of data points required to come up with a robust spreadsheet becomes impossible. Our expertise in machine learning and neural networks were called upon to solve this problem. We were able to improve upon their existing method under various evaluation metrics. During this project, the nature of this data also inspired some interesting changes to the classic neural network architecture, which are currently under exploration.

**Kuan-Chieh Wang**  
kc.wang@utoronto.ca

**Prof. Richard Zemel**  
zemel@cs.toronto.edu

---

### Real-time Coarse to Fine Topologically Preserving Segmentation | BOOTH 16 (AI)



In this poster, we tackle the problem of unsupervised image segmentation in the form of superpixels. Our main emphasis is on speed and accuracy. We define the problem as a boundary and topology preserving Markov random field. We propose a coarse to fine optimization technique that speeds up inference in terms of the number of updates by an order of magnitude. Our approach is shown to have promising results while employing a single iteration. We evaluate and compare

our approach to state-of-the-art superpixel algorithms on the BSD and KITTI benchmarks. Our approach significantly outperforms the baselines in the segmentation metrics and achieves the lowest error on the stereo task.

**Jian Yao**  
zyyaojian@gmail.com

**Prof. Raquel Urtasun**  
urtasun@cs.toronto.edu

---

### Ranking Networks | BOOTH 41 (AI)




Latent space models for network formation assume that nodes possess latent attributes that determine their propensity to connect. We propose a new model for network formation, ranking networks, in which these attributes are rankings over some space of alternatives. Such rankings may reflect user preferences, relevance/quality judgements, etc., while ranking networks capture correlations of, say, user preferences across a social network.

We present preliminary theoretical and empirical analyses of structural properties of such networks, and develop algorithmic approximations to help efficiently predict these properties. Empirical results demonstrate the efficiency and accuracy of quality of these approximations.

**Amirali Salehi-Abari**  
abari@cs.toronto.edu

**Prof. Craig Boutilier**  
cebly@cs.toronto.edu


## Stochastic Segmentation Trees | BOOTH 43 (AI)



Many structured output problems such as image segmentation admit multiple correct outputs for a single input. We present an agglomerative approach for modeling multiple output segmentations via a hierarchical tree of image regions. We perform learning by matching merge probabilities of the model with those observed in the ground truth distribution. The hierarchical nature of the resulting model allows for fast inference and sampling. We conduct experiments on segmentations from both the Penn-Fudan pedestrian dataset and the Berkeley Segmentation dataset.

**Jake Snell**  
jsnell@cs.toronto.edu

**Prof. Richard S. Zemel**  
zemel@cs.toronto.edu



## Writing Systems | BOOTH 28 (AI)

The decipherment of Egyptian hieroglyphics, Linear B, and Mesopotamian Cuneiform have all inspired research in writing systems. It includes work on classifying the different types of writing systems, the direction of reading, the grammatical structure of a text, the sound values associated to the writing system's symbols, and so on. A non-exhaustive list of different types of writing systems includes logographic systems (e.g., Chinese), syllabic systems (e.g., Japanese kana), and alphabetic systems (e.g., the one used to write this blurb). Research results in this area are often put into practice in lost language decipherment, machine translation, text-to-speech applications, and analysis of language change through time. The project at hand uses problems faced in lost language decipherment as a road map for describing the state-of-the-art approaches in writing systems research.

**Patricia Thaine**  
pthaine@cs.toronto.edu

**Prof. Gerald Penn**  
gpenn@cs.toronto.edu

## ARTIFICIAL INTELLIGENCE (AI) | MOBILE COMPUTING (Mobile)

### Data Mining and Machine Learning in Mobile App Advertising | BOOTH 27 (AI + Mobile)



Addictive Mobility is a leading Canadian mobile advertising company providing a technology platform for advertisers, brands, etc to have their ads shown in a large inventory of mobile apps. User engagement and performance of ads are the selling point of the company's products and optimizing these metrics is the goal of the Data Science team. We try to optimize performance using different machine learning approaches that mainly predict certain aspects of an ad campaign or attributes about our users.

Our research internship involved the inception and development of a user profiling system, which analyzes the massive amount of data that we receive and extract useful user profiles. A system to predict the age and gender of a user based on their profiles was also developed. A current project is segmentation and clustering of users based on their behavior and interests. We also built a low latency decision module integrated with Addictive Mobility's real-time bidding system, which makes decisions on which ad impressions to bid on. Some of the factors it incorporates includes a statistical model predicting user engagement level, and a smart bid price computation tool based on historical bid data.

**Lubna Khader, Megha Lakshmi Narayanan and Guilherme M. Trein**

lubnakhader@cs.toronto.edu

**Prof. Anthony Bonner**

bonner@cs.toronto.edu

---

## COMPUTATIONAL BIOLOGY (CompBio)

### A Tale of Two Parasites: The Central Role of Energy Metabolism in *Sarcocystis neurona* and *Toxoplasma gondii* | BOOTH 47 (CompBio)


The closely-related parasites *Sarcocystis neurona* and *Toxoplasma gondii* have spread with great success, having evolved the ability to infect different tissue types and hosts. Both are highly medically-relevant: for example, *T. gondii* is the leading cause of childhood retinitis. Due to the emergence of drug-resistant strains, new treatments are urgently required. Here, we focus on understanding how metabolic potential can drive the virulence of these parasites. We apply Flux Balance Analysis (an extension of Linear Programming) augmented with gene expression data to understand what drives the virulence of these parasites. In the case of *T. gondii*, we observe that different strains rely on different systems/pathways. In fact, alternate usage of energy-producing pathways results in different growth rates in different strains of the parasite—a result verified through a drug inhibition assay. We have then extended our work to *S. neurona*. We predict that *T. gondii* and *S. neurona* rely on different sets of pathways for growth. More interestingly, in *S. neurona*, we predict that different energy-producing pathways are used depending on the availability of either sucrose or glucose in the environment—indicating a degree of robustness in its metabolic network, as well as hinting at the key to the parasite's successful spread to different hosts. Overall, these studies showcase the power of FBA in elucidating the central role of energy metabolism in parasites.

**Nirvana Nursimulu**

nirvana.nursimulu@mail.utoronto.ca

**Prof. John Parkinson**

john.parkinson@utoronto.ca



**Expanding the Epigenetic Alphabet: A New Dimension in Transcription Factor Binding Prediction | BOOTH 52 (CompBio)**

Many transcription factors show affinities for particular DNA sequences, providing the means for sequence specificity of transcriptional control. The position weight matrix (PWM) model allows for the computational identification of transcription factor binding sites (TFBSs), by characterizing a transcription factor's position-specific preference over the DNA alphabet. This four-letter alphabet, however, only partially describes the possible diversity of nucleobases a transcription factor might encounter. Cytosine is often present in a covalently modified form: 5-methylcytosine (5mC). It can be further modified to form: 5-hydroxymethylcytosine (5hmC), 5-formylcytosine (5fC), and 5-carboxylcytosine (5caC). Just as transcription factors distinguish one unmodified nucleobase from another, some transcription factors are known to distinguish between unmodified and modified bases, and transcription factors are known to often bind to novel modified motifs. These newer epigenetic DNA modifications' effects on transcription factor binding remain uncharacterized, as the basic tools for their analysis do not yet exist. To understand the effect of modified nucleobases on gene regulation, we developed methods to discover transcription factor motifs and identify TFBSs in DNA with covalent modifications. Our models expand the standard A/C/G/T alphabet, adding m (5mC), h (5hmC), f (5fC), and c (5caC). We have engineered several tools to work with this expanded-alphabet sequence and PWMs. These tools permit us to integrate epigenomic assays to create a modified genome sequence and to perform *de novo* motif elucidation and characterization over this expanded alphabet. Using these methods, we have been able to identify *cis*-regulatory modules that we believe are either active or inactive contingent upon the presence or absence of cytosine modifications.

**Coby Viner**  
 cviner@cs.toronto.edu


**Prof. Michael M. Hoffman**  
 michael.hoffman@utoronto.ca

**Protein Domain-level Landscape of Cancer-Type-Specific Somatic Mutations | BOOTH 39 (CompBio)**

Identification of driver mutations and their functional consequences is critical to understand cancer processes and design novel therapeutics. Towards this goal, and because domains are the functional units of a protein, we explored the protein domain-level landscape of cancer-type-specific somatic mutations across 21 cancer types. Within domains, we systematically identified the position, as well as the functional and structural contexts of mutational hotspots in different cancer types. While hotspots corresponding to specific gain-of-function mutations are expected for oncoproteins, we found that tumor suppressor proteins also exhibit strong biases toward being mutated in particular domains. Within domains, however, we observed the expected pattern of mutation for oncogenes and tumor suppressors, with recurrently mutated positions and evenly distributed mutations, respectively. For example, we identified both known and new endometrial cancer hotspots in the tyrosine kinase domain of the FGFR2 protein, one of which is also a hotspot in breast cancer, and found new two hotspots in the Immunoglobulin I-set domain in colon cancer. Thus, to prioritize cancer mutations for further functional studies aimed at more precise cancer treatments, we have systematically correlated mutations and cancer types at the protein domain level.

**Fan Yang**  
 sunnyfan.yang@mail.utoronto.ca

**Prof. Frederick Roth**  
 fritz.roth@utoronto.ca



**PULSE: Positive Unlabelled Learning for Splicing Elucidation | BOOTH 40 (CompBio)**

Alternative splicing acts on transcripts from almost all human multi-exon genes. Notwithstanding its ubiquity, fundamental ramifications of splicing on protein expression remain unresolved. The number and identity of spliced transcripts that form stably folded proteins remains the source of considerable debate, largely due to low coverage of experimental methods and the resulting absence of negative data. We circumvent this issue by developing a semi-supervised learning algorithm, Positive Unlabelled Learning for Splicing Elucidation (PULSE), which uses 48 features spanning various categories. We validated its accuracy on sets of bona fide protein isoforms and directly on MS-spectra for an overall AU-ROC of 0.85. We predict that around 32% of ‘exon skipping’ alternatively splicing events produce stable proteins, suggesting that the process engenders a significant number of novel proteins. We also provide new insights into the distribution of positive isoforms in various functional classes, and into the structural effects of alternative splicing.

for Splicing Elucidation (PULSE), which uses 48 features spanning various categories. We validated its accuracy on sets of bona fide protein isoforms and directly on MS-spectra for an overall AU-ROC of 0.85. We predict that around 32% of ‘exon skipping’ alternatively splicing events produce stable proteins, suggesting that the process engenders a significant number of novel proteins. We also provide new insights into the distribution of positive isoforms in various functional classes, and into the structural effects of alternative splicing.

**Recep Colak**  
rcolak@cs.toronto.edu

**Prof. Philip M. Kim**  
pm.kim@utoronto.ca

**GAMES**



**Afterlight | BOOTH 33 (Game)**

Afterlight is a single-player puzzle platformer video game made in CSC404 by Third Eye Studio- a team of 7 U of T and OCAD students. The game takes place in the monochromatic dual worlds of Invertia, where the player must use their ability to switch between the worlds to solve puzzles and head towards the top of the great Clock Tower. Afterlight won the *People’s Choice Award* at the Level Up 2015 Showcase.

**Cordelia Leung**  
cordelia.leung@mail.utoronto.ca

**Spencer Christie**  
spencer.christie@mail.utoronto.ca

**Tony Hong**  
tonyxu.hong@gmail.com

**Steve Engels**, Senior Lecturer  
sengels@cs.toronto.edu

**Saffron Bolduc-Chiong, Prisilla Hui, Lauren Thomas and Nick Wang**, OCAD U

**Annexing Reality | BOOTH 15 (Game)**

Annexing reality is the use of everyday physical objects as tangible proxies for virtual content in an Augmented Reality context.

**Anuruddha Hettiarachchi**  
anuruddha@dgp.toronto.edu

**Prof. Daniel Wigdor**  
daniel@dgp.toronto.edu





### **Avoid | BOOTH 32 (Game)**

A\*VOID is a video game of space sumo dodgeball. In this game, two to four players battle it out in a top-down perspective space arena. Players combat one another by shooting stars at each other, and by forming strategically placed voids (aka black holes). Only one player can win by being the last person standing. A\*VOID was showcased at the Level Up 2015 Student Games Showcase. A\*VOID was a collaboration between University of Toronto and

OCAD students. A\*VOID was built using Unity3D and C#.

#### **Jiang Tao Hong**

jiang.hong@mail.utoronto.ca

#### **Eric Jimenez**

eric.jimenez@me.com

#### **Michael Wehrmann-Kaminsky**

mwk112358@gmail.com.

#### **Zhi Zhang**

zhi.study.ut@gmail.com

#### **Steve Engels, Senior Lecturer**

sengels@cs.toronto.edu

#### **Lionel Grant, Jessie Sheng and Ashley Yip, OCAD U**

---

### **Game-Changing Features for Blind Players | BOOTH 30 (Game)**

The game design process integrates visual, audio, gameplay, and control elements into a single experience. Of these, the audio component is often treated as a secondary concern in both the design and research arenas. Our research seeks to remedy this, by focusing on audio elements of games, the impact they can have, and how audio design can be integrated in the overall design process. With the rising popularity of game development and gaming, there exists a concern for users who have low vision and blindness. This challenge has prompted research on the development of games for the blind comprised of both people with and without blindness developers. Our poster will present the findings from our research on designing accessible games for the blind. We will discuss design guidelines and recommendations in developing content for this audience.

#### **Tiffany Tong**

tiffany.tong@mail.utoronto.ca

#### **Steve Engels, Senior Lecturer**

sengels@cs.toronto.edu

#### **Dr. Daniel Zingaro, Lecturer**

daniel.zingaro@utoronto.ca



### **Pitfall Planet | BOOTH 31 (Game)**

Pitfall Planet, our project for CSC404, is a cooperative puzzle-solving game. It takes place in an xtra-terrestrial mine, where two greedy astronauts must collaborate to collect ore and make their way down deeper into the mine. The levels consist of puzzles that are best solved when the players work together. Each astronaut is equipped with a grappling hook that he or she can use to interact with objects in the environment. Pitfall Planet won *Best Overall*

*Game* at the 2015 Level Up Showcase and *2nd Best Art Design*, competing against 82 other teams.

#### **Alexander Biggs**

alexanderbiggs@gmail.com

#### **Daphne Ippolito**

daphneipp@gmail.com

#### **Adam Robinson-Yu**

adamgryu@gmail.com

#### **Steve Engels, Senior Lecturer**

sengels@cs.toronto.edu

#### **David Czarnowski, Emma Marie-Burkeitt and Guy Torcher**

OCAD U

## HUMAN-COMPUTER INTERACTION/GRAPHICS (HCIG)

### FlatFab and Janus VR | BOOTH 23 (HCIG)



flatfab is software that helps you design and fabricate 3D objects. We imagine in 3D but ideate and draw in 2D. flatfab seamlessly does both. ...awaken the dormant designer in you. Janus VR allows a spatial walk through the internet, inspired by the novel Snow Crash

by Neal Stephenson who detailed a metaverse. Our analogy is that webpages are rooms, and links connect rooms via portals (doorways which seamlessly connect rooms). Try out the latest virtual reality technology and join us in the web!

**James McCrae**  
mccrae@dgp.toronto.edu

**Prof. Karan Singh**  
karan@dgp.toronto.edu

---

### A Transport-Aware Video Camera | BOOTH 55 (HCIG)




(a) conventional image



(b) indirect image



(c) direct image



(d) our primal-dual coding cam

Conventional cameras record all light falling on their sensor regardless of the path that light followed to get there. We will demonstrate a new type of video camera that offers many more degrees of freedom. This camera records just a fraction of the light

coming from a controllable source, based on the actual 3D light path followed. Photos and live video captured this way offer an unconventional view of everyday scenes in which the effects of scattering, refraction and other phenomena can be selectively blocked or enhanced, visual structures that are too subtle to notice with the naked eye can become apparent, and object appearance can depend on depth. We will show how to use the camera to visualize indirect transport effects (caustics, inter-reflections, volumetric scattering, etc.), measure direct-only light paths (surface reflections), capture transport that occurs within specific regions of 3D space, and perform structured-light 3D imaging in the presence of complex indirect effects.

**John Mather**  
jmather@dgp.toronto.edu

**Prof. Kyros Kutulakos**  
kyros@cs.toronto.edu

**Matthew P. O'Toole**  
motoole@cs.toronto.edu


### **Reach – Grip Sensing for Interactions in Large Screen Mobile Devices | BOOTH 10 (HCIG)**

The benefits of smart phones with large screens drive their popularity. However these devices are difficult, if not impossible, to use with one hand and pose usability issues for demographics with smaller hands. The existing solutions to this include on-screen functions that the user can activate to bring the screen content closer to the thumb. These methods however introduce extra steps in the users' interaction with the device and can be cumbersome.

By placing force sensors all around the rim of the phone, we can sense how the user is holding the phone and when they are straining their thumb to reach a corner. Using this information we can shift the UI closer to the operating finger. The force sensors can also be used to interact with the phone in other scenarios, for example swiping on the sides of the phone could scroll pages, or increase/decrease volume etc. With this project we intend to build the hardware, formulate UI design changes, and do basic user testing to validate our ideas.

**Varun Perumal Chadalavada**  
varun@dgp.toronto.edu

**Prof. Daniel Wigdor**  
daniel@dgp.toronto.edu



### **An Interactive Serious Game For Cognitive Screening | BOOTH 48 (HCIG)**

The demographic transition to increasingly older populations in developed countries is likely to create an explosion in age-related conditions such as mild cognitive impairment and dementia. Healthcare systems need to adapt to this challenge and one necessary adaptation involves more efficient and pervasive screening for cognitive health. Current cognitive testing methods in the elderly rely on clinical assessments, which are time consuming, costly, and require highly trained staff (Kueider et al. 2012). Assessing mental health is an important indicator of functional ability, independence, and mortality (Inouye et al. 2000). The under-diagnosis of cognitive impairments can lead to an increased economic burden, hospitalization, and even death (Inouye et al. 2000). Many of the current tests have been developed to screen for specific conditions, but there is a lack of cognitive assessments to screen for transitory conditions that occur between normal cognition and cognitive failure. This highlights the need for a patient-administered test that monitors transitions towards break down of cognition. To address the challenge, we have developed a serious game that assesses cognition, which has undergone several usability studies and is currently undergoing a validation study in a clinical setting. Our findings reveal that our game-based assessment strongly correlates with the executive function of inhibition, which is subject to age-related decline.

**Tiffany Tong**  
tiffany.tong@mail.utoronto.ca

**Prof. Mark Chignell**  
chignell@mie.utoronto.ca

**Steve Engels**, Senior Lecturer  
sengels@cs.toronto.edu

### **Device Notation: Composing Cross-Device Experiences Using Physical Tags | BOOTH 9 (HCIG)**

Our world is quickly becoming permeated with interconnected devices. However, experiences that leverage the capabilities of multiple devices are rare. This is largely because these devices are designed for independent use and there is no mechanism that allows users to define inter-device relationships. In our project, we attach touch-activated tags to devices. The tags contain identifying information read by a wrist-worn sensor. This information can be transformed by a user's personal devices (i.e. phone, smart watch) into contextually aware interfaces for the tagged device. The information can also be used to facilitate the formation of connections between devices (e.g. keyboard and tablet).

**Peter Hamilton**  
hamilton@dgp.toronto.edu

**Prof. Daniel Wigdor**  
daniel@dgp.toronto.edu



## Direction Manipulation of Attributes in Authoring Tools | BOOTH 18 (HCIG)

Direct manipulation has become the dominant approach of interaction with graphical user interfaces since the mouse was invented: we can move the cursor to any location on the screen, select an object, and drag it around. The recent proliferation of touch screens enables a more direct way: we can interact with objects directly with our hands. Unfortunately, what has not changed is how we manipulate the attributes of an object. Unlike an object, attributes are virtual values that define an object's appearance and its relationship to other objects. Currently we have to access the attributes through buttons and menus. This brings additional difficulty to the interaction in that the interface gets very complex. In this design, attributes are no longer values buried underneath buttons and menus. Attributes are objects. By giving embodiment to attributes, we can leverage our experience of interaction with physical objects in the real world, enabling many new interaction possibilities.

**Haijun Xia**

haijunxia@dgp.toronto.edu

**Prof. Daniel Wigdor**

daniel@dgp.toronto.edu

---

## Ghost trick: Selecting and Animating Skinned Mesh with Mimic Gesture | BOOTH 22 (HCIG)



Tired of controlling a human character to explore the virtual world? Now, use your gesture to mimic the motion of a 'character' (e.g. animals, monsters) in your mind, and our technique will assist you in selecting and animating a skinned mesh in the virtual world.

**Yupeng Zhang**

yupeng@dgp.toronto.edu

**Prof. Daniel Wigdor**

daniel@dgp.toronto.edu

---

## Storeboard: Sketching Stereoscopic Storyboards | BOOTH 20 (HCIG)



For this project, we designed a prototype to explore creating quick stereoscopic storyboards in real-time. The interface is modeled after a sketchpad and pencil, removing many GUI elements (such as dropdown lists and menus) and negates the necessity for environmental controls (such as cameras and viewport widgets). The program explores numerous interaction techniques that leverage different combinations of pen and touch input, and provides tools to manipulate objects, layers, and planes. Finally we examined unique problems associated with working in a planed stereoscopic 3D environment, and investigated potential solutions to these problems.

**Rorik Henrikson**

rorik@dgp.utoronto.ca

**Prof. Ravin Balakrishnan**

ravin@dgp.utoronto.ca

**Prof. Karan Singh**

karan@cs.toronto.edu

**Dr. Bruno De Araujo, PDF**

brar@cs.toronto.edu



### PhenoBlocks: Phenotype Comparison Visualizations | BOOTH 19 (HCIG + CompBio)

The differential diagnosis of hereditary disorders is a challenging task for clinicians due to the heterogeneity of phenotypes observed in patients. Existing clinical tools are often text-based and do not emphasize consistency, completeness, or granularity of phenotype reporting, which can impede clinical diagnosis and can limit downstream utility to genetics researchers.

We present PhenoBlocks, a novel visual analytics tool that supports the pairwise comparison of phenotypes between patients. Not only can PhenoBlocks aid clinicians in building complete and detailed phenotypic descriptions of patients, but also functions as a common visual language between clinicians and researchers, helping to improve future research outcomes.

**Michael Glueck**  
mglueck@dgp.toronto.edu

**Peter Hamilton**  
hamilton@dgp.toronto.edu

Prof. **Daniel Widgor**  
daniel@dgp.toronto.edu

Prof. **Michael Brudno**  
brudno@cs.toronto.edu

**Fanny Chevalier, Inria Simon Breslav and Azam Khan**  
Autodesk Research

---

## MOBILE

### Home Page Layout Optimization | BOOTH 24 (Mobile)

Kobo is an eBook company dedicated to providing a world-class eReading platform to millions of users in 190 countries. Kobo offers one of the world's largest catalogues. With over 5 million books to browse, content discovery can be a daunting task for readers. In order to optimize content discovery, we adapt our website content to targeted groups of readers by making use of their purchasing, browsing and demographic data. We model targeted group content recommendations as a contextual bandit problem, in an exploration-exploitation scheme, in which an algorithm dynamically selects the optimal website layout for a pool of users, while simultaneously adapting the content of said website layout to maximize conversion.


**Sagun Bajracharya**  
sbajra@kobo.com

Prof. **Brendan Frey**  
frey@psi.toronto.edu

**Inmar Givoni**, VP of Big Data  
Kobo

---

### Innovative Spiral Lock Design for Smartphone Security | BOOTH 7 (Mobile)



We introduce a new locking concept to improve setbacks of Pattern and PIN locks to find leads towards a more convenient and secure unlocking gesture. We hypothesize our design will outperform the Pattern lock in both efficiency and failed unlockings, and believe the Pattern and our design will underperform in comparison to the PIN lock. In our study with twenty-seven participants, our design performed better in comparison to the average speed of the Pattern and PIN lock, although the Pattern performed better than the PIN, probably because of sample size. In terms of errors,

there were no significant differences between Pattern and our design, but as suspected the PIN lock performed better than both in this regard. Unfortunately, we received qualitative feedback that our design is the least comfortable amongst the three locks, but is of preference similar to Pattern.

**Jacqueline Bermudez, Benson Quach, Alaa Abdulaal and Liwen Xu**  
jacqueline.bermudez@mail.utoronto.ca

## MOBILE | HUMAN-COMPUTER INTERACTION/GRAPHICS

### Sharing Application Usage in Public Spaces | **BOOTH 21** (Mobile + HCI/G)



Double-sided tablets  
Sharing Application Usage in Public Spaces

Our suite of mobile devices are inherently private, even when we use them in a public space. People lament that they text others or direct their attention to technology instead of focusing on the people around them, whether they're at a bar, having dinner or at work. This project explores what happens when users share the applications they're currently using on phones, tablets and laptops with people around them. How does this additional information impact the public space users are in, their use of the technology

and the in-person interactions they have with others?

We have two main studies. The first explores an open office environment where people are constantly sharing their application usage to public screens displayed throughout the office. The second study explores a number of prototypes which introduce a second public display for every private display we have. Example prototypes include: double sided tablets which have a public screen on the back in addition to private screen on the front, or phones that project what we're doing on the ground in front of us.

**Adam Hueniken**  
adam.hueniken@mail.utoronto.ca

Prof. **Daniel Wigdor**  
daniel@dgp.toronto.edu

---

## NUMERICAL ANALYSIS (NA)


### New Approaches to Importance Sampling for Portfolio Credit Risk Valuation | **BOOTH 35** (NA)

Portfolio credit risk based on the Gaussian Copula model has been widely studied and generally evaluated through Monte Carlo simulations. The two-level structure, namely systematic factors and individual factors, complicates the problem in a way that makes traditional variance reduction techniques hard to apply. Glasserman and Li proposed a two-level importance sampling approach to tackle a simplified binary-credit-states problem. The inner level is approximated by a conditional importance sampling approach using an exponential twisting technique. We propose an alternative importance sampling approach which uses the Central Limit Theorem for the inner level. Our approach can be easily generalized to multi-credit states. Based on this approximation, we then propose two novel approaches motivated from research in machine learning. Instead of finding the importance sampler through an optimization problem, the first approach approximates the zero variance function by learning from the samples which are generated from Markov Chain Monte Carlo. The second approach treats the problem as a Bayesian inference problem and evaluates the tail probability through Bayesian Monte Carlo. Compared to Glasserman and Li's method, numerical results show that these two new approaches have advantages in both accuracy and speed. Moreover, they are also more easily adapted to other applications.

**Zhe (Robert) Wang**  
zwang@cs.toronto.edu

Prof. **Ken Jackson**  
krj@cs.toronto.edu

### A 3-stage Parameter Estimation Procedure for ODEs and DDEs | [BOOTH 45 \(NA\)](#)



We propose a 3-stage procedure for performing parameter estimation for ODE and DDE models. The first stage attempts to obtain improved initial guesses for a subset of the model parameters. The second stage uses a global optimizer to get us 'close' to the minimum value of our least squares objective function. The third stage uses a local optimizer to obtain the parameter estimates. In this work, we considered the case where measurements are available for the entire state vector, at specified values of the independent variable. We report the cost of


performing the parameter estimation for a predator-prey model and show how each stage contributes to reducing the total cost.

**Jonathan Calver**  
calver@cs.toronto.edu

**Prof. Wayne Enright**  
enright@cs.toronto.edu

---

## PEDAGOGY



### Courseography | [BOOTH 26 \(Pedagogy\)](#)

Courseography is a web application designed to help students plan their courses at the University of Toronto. The two main components are an interactive course graph displaying prerequisite relationships in the Department of Computer Science, and a timetable grid allowing students to quickly build their schedules for the current academic year, drawing from courses across the Faculty of Arts & Science. This year, students have worked on

integrating program requirements, developing graph generation tools to facilitate the entry of other departments to Courseography, and exporting data to both static file formats and Facebook.

**Ian Stewart-Binks**  
ian.stewart.binks@mail.utoronto.ca

**David Liu, Lecturer**  
liudavid@cs.toronto.edu

**Paul Gries, Sr. Lecturer**  
pgries@cs.toronto.edu

---

### MyLiveNotes | [BOOTH 13 \(Pedagogy\)](#)

myLiveNotes by Crescendum is a web-based system designed with academic reading in mind. With an intuitive interface that scales to any device's screen size, anybody can easily study, annotate and share online content. The goal for myLiveNotes is to create a tool that empowers educators and learners to share diverse formats of relevant content that can be actively accessed and understood, worked upon within a group setting, and shared with peers. The system is powered by Ruby on Rails and React using the Flux application architecture.

**Jonathan Webb**  
jonathan.webb@mail.utoronto.ca

**Paul Gries, Senior Lecturer**  
pgries@cs.toronto.edu

### **Automatic Assessment and Feedback for Online C Programming Exercises | BOOTH 11 (Pedagogy)**

Recently, best practice in computer science education has shifted to “inverted classrooms” where students are introduced to and practice skills outside of the classroom, delve into the content in depth in the classroom, and then rehearse the skills developed. For this structure to function in a programming course, instructors must be able to craft exercises that target very specific content areas, and students must be provided with opportunities to write code and to receive timely, accurate feedback on their solutions.

The Programming Course Resource System (PCRS) provides instructors with the ability to create online modules, consisting of videos, text, and a variety of exercises, including programming exercises. Recently, PCRS was extended to support the C programming language. The system allows instructors to define programming problems which focus on specific components of C programs (hiding irrelevant or unimportant code). Students receive instantaneous feedback on their submissions, based on a test suite provided by the instructor, and they can also use a memory visualizer to trace through program execution.

**Daniel Parreira**

daniel.parreira@udoctor.com.br

**Andrew Petersen**, Senior Lecturer (UTM)

andrew.petersen@utoronto.ca

**Michelle Craig**, Senior Lecturer

mcraig@cs.toronto.edu



### **MarkUs | BOOTH 25 (Pedagogy)**



MarkUs is a web application for the submission and grading of student programming assignments. The primary purpose of MarkUs is to provide TAs with simple tools that will help them to give high quality feedback to students. MarkUs also provides a straight-forward interface for students to submit their work, form groups, and receive feedback. The administrative interface allows instructors to manage groups, organize the grading, and release grades to students.

Since 2008, more than 140 undergraduate students have participated in the development of MarkUs; some as full-time summer interns, but most working part time on MarkUs as a project course. The fact that we have uncovered so few major bugs, and that MarkUs has been so well-received by instructors is a testament to the high quality work of these students. MarkUs is used in more than a dozen courses at the University of Toronto, in several courses at the University of Waterloo, and at École Centrale Nantes (in French).

**Karen Reid**, Senior Lecturer

reid@cs.toronto.edu

**David Liu**, Lecturer

liudavid@cs.toronto.edu



## **RAPT | BOOTH 12 (Pedagogy)**

Many database courses rely on relational algebra (RA) to provide a theoretical foundation for database query languages such as SQL. However, few tools exist to support students in learning RA. To fill this need, we created RAPT. RAPT uses a syntactic and semantic understanding of RA to transform input statements into a variety of outputs, including latex formatted queries, parse tree diagrams, and executable SQL statements. The translation to SQL is particularly important, as it enables the creation of automatically tested exercises and allows students to view the result of an RA statement. RAPT-supported exercises have been integrated into PCRS, a system for creating online learning modules and deployed to a third year databases course with over 350 students.

**Noel D'Souza**  
noel.dsouza@mail.utoronto.ca

**Diane Horton**, Senior Lecturer  
dianeh@cs.utoronto.ca

**Olessia Karpova**  
olessia.karpova@mail.utoronto.ca

**Andrew Petersen**, Senior Lecturer (UTM)  
andrew.petersen@utoronto.ca

---

## **Persistence and Outcomes in Inverted and Traditional CS1 | BOOTH 37 (Pedagogy)**

Much attention has been paid to the failure rate in CS1 and attrition between CS1 and CS2. In our study of 1236 CS1 students, we examine subgroups of students, to find out how characteristics such as prior experience and reason for taking the course influence who drops, fails, or passes, and who continues on to CS2. We also examine whether student characteristics influence outcomes differently in traditional vs inverted offerings of the course. We find that more students in the inverted offering failed the midterm test, but those who failed were much more likely to either drop the course or recover and ultimately pass the course. While we find no difference between the offerings in the overall drop-fail-pass rates or in the percentage and types of students who go on to take CS2, there is a significant, widely felt, boost in exam grades in the inverted offering.

**Diane Horton**, Senior Lecturer  
dianeh@cs.utoronto.ca

**Jennifer Campbell**, Senior Lecturer  
campbell@cs.toronto.edu

**Michelle Craig**, Senior Lecturer  
mcraig@cs.toronto.edu

**Paul Gries**, Senior Lecturer  
pgries@cs.toronto.edu

**Dr. Daniel Zingaro**, Lecturer (UTM)  
daniel.zingaro@utoronto.ca

---

## **SOFTWARE ENGINEERING (SE)**


### **Analyzing IT Flexibilities in Enterprises | BOOTH 6 (SE)**

Enterprise Information Systems are moving from monolithic applications to loosely coupled services including cloud and other networked architectures. These architectural approaches tend to reduce upfront effort but can potentially create integration burdens down the road. If the ultimate objective is enterprise-level flexibility, how effective are lightweight orchestrated software applications for achieving this flexibility? In this research we propose a conceptual modeling framework that can answer questions regarding how to develop, orchestrate and deploy IT capabilities and systems in accordance with organizational characteristics and requirements. We demonstrate the framework's ability to facilitate decision making regarding a choice between RESTful and centralized ESB implementation for business process management. The systemic analysis not only depicts the tradeoffs among the choices but also illustrates the impact on organizational flexibility in each case.

**Mohammad Hossein Danesh**  
danesh@cs.toronto.edu

**Prof. Eric Yu**  
eric@cs.toronto.edu

### City on the River: Visualizing Individual & Group Collaboration Over Time | **BOOTH 2 (SE)**



In the recent decade, to solve large and complicated problems, people are placed into teams with the intent that they'll use their different abilities to meet the project requirements and share the workload. In these collaborative group settings, it is interesting to be able to analyze the collaboration over time, individual contributions, and group performance. To help accomplish this

task, we developed a visualization approach that maps the collaborators to their products over time along with their in-team and out-team attributes. In addition to three usage scenarios, our demo will include a prototype with two synthetic data sets based on a student software engineering class project and a researcher network program. We will also describe our proposed study plan for evaluating the effectiveness of the visualization.

**Jessica Perrie**  
perrie@cs.toronto.edu

**Prof. Kelly Lyons**  
kelly.lyons@utoronto.ca

---

### MU-MMINT: an IDE for Model Uncertainty | **BOOTH 14 (SE)**



Developers have to work with ever-present design time uncertainty, i.e., uncertainty about selecting among alternative design decisions. However, existing tools do not support working in the presence of uncertainty, forcing developers to either make provisional, premature decisions, or avoid using the tools altogether until uncertainty is resolved. We present a tool, called MU-MMINT,

that allows developers to express their uncertainty within software artifacts and perform a variety of model management tasks such as reasoning, transformation and refinement in an interactive environment. This allows developers in turn to defer the resolution of uncertainty, thus avoiding having to undo provisional decisions. See also the companion video: <http://youtu.be/kAWUm-iFatM>

**Michalis Famelis, Naama Ben-David, Alessio Di Sandro and Rick Salay**  
famelis@cs.toronto.edu

**Prof. Marsha Chechik**  
chchik@cs.toronto.edu

---

### Engaging End-Users in Green Building Design Software | **BOOTH 51 (SE)**

Green building design is a socio-technical process in which end-users should be engaged from the beginning. Given that a majority of software tools used in designing green buildings are aimed at engineers, we must determine the most effective way of communicating information about energy use to end-users, who are typically unfamiliar with energy analysis techniques. Research has not yet determined how best to present environmental impact information to end-users in order to engage them in the design process. In this project, we will build software that brings professionals and end-users together in a social media platform and enables socio-technical analytics of green buildings. Through an ethnographic study we hope to determine an effective method of presenting information to end-users within this platform. The project's goal is to get end-users invested in the environmental performance of their building and provide them with the knowledge and tools to become conscious energy consumers.

**Mahtab Sabet**  
mahtab@cs.toronto.edu

**Prof. Steve Easterbrook**  
sme@cs.toronto.edu

## Data Analytics for Adaptive Enterprise | BOOTH 5 (SE)



Increasingly complex and uncertain business environments, rapid evolution of data, as well as the need to sense and respond to market trends and changes have made data analytics a strategic priority of organizations. Nonetheless, to effectively implement data analytics and exploit its potential is a non-trivial task. A critical challenge in these contexts is to make data analytics and business intelligence systems more actionable, which includes systematically connecting and coordinating analytics-driven

insights with consequent enterprise actions, aligning them with business strategies and users' requirements, and really using them to improve the business situation. This research addresses these problems by developing, extending, and using requirements engineering and conceptual modelling frameworks to make a connection between what an enterprise senses from BI-driven insights and consequent actions and corrective changes. The proposed framework (i) connects analytics-driven insights to enterprise actions; (ii) facilitates decision making over discovered insights, which in-turn will assist exploration of the next-best action; and (iii) comes with a set of design patterns to support modeling and reasoning about data analytics applications in enterprise architectures. Case studies show that this framework can support design and use of data analytics systems, connecting their insights to actions, hence facilitating enterprise adaptiveness.

**Soroosh Nalchigar**


soroosh@cs.toronto.edu

**Prof. Eric Yu**

eric@cs.toronto.edu

---

## S+EPP Scalable Graph Summary Construction and Navigational Query Optimization | BOOTH 29 (SE)



We demonstrate S+EPP, a system that provides fast construction of bisimulation summaries using graph analytics platforms, and then enhances existing SPARQL engines to support summary-based exploration and navigational query optimization. The construction

component adds a novel optimization to a parallel bisimulation algorithm implemented on a multi-core graph processing framework, achieving the goal of constructing summaries in roughly the same time as the data load (for several large, disk resident, real world graphs). The query component supports EPPs (Extended Property Paths, an enhancement of SPARQL 1.1 property paths that can express a significantly larger class of navigational queries) by implementing a translation from EPPs to a widely implemented SPARQL subset. The optimization component can (transparently to users) translate EPPs defined on instance graphs into EPPs that take advantage of the summary graphs created by S+EPP. Combining both translations enables summary-based optimization of graph traversal queries on top of SPARQL processors. The demonstration showcases S+EPP's ability to construct bisimulation summaries of graph datasets (ranging from millions to billions of edges), together with the exploration benefits and the navigational query speedups obtained by leveraging summaries stored alongside the original datasets.

**Shahan Khatchadourian**

shahan@cs.toronto.edu

**Prof. Mariano P. Consens**

consens@cs.toronto.edu

**Valeria Fionda**

University of Calabria

**Giuseppe Pirro**, CNR

## Designing Software Ecosystems: Methods, Techniques and Tools | BOOTH 44 (SE)



The software industry is in transition to a new era, where there is rapid technical innovation. To accommodate the dynamicity of the new environment, many software companies are adopting open architectures, open innovation models, and collaborative software development practices, constituting software ecosystems. This paradigm shift demands new thinking in the architecting of software products and services, such as the difficult balancing of openness and community building with proprietary product development and revenue generation. In this new environment,


software engineering has become entwined with strategic business decision making, and the orchestration of social relationships among many collaborating partners, internal and external, some of whom are also competitors in the marketplace. In this research, we develop methods and tools, as well as the underlying conceptual foundations, to support the ongoing operation and the longer-term viability and sustainability of ecosystem software initiatives. We identify critical factors for attracting participants, for accelerating adoption, and for sustaining the longer-term viability of an ecosystem.

**Mahsa H. Sadi**  
mhsadi@cs.toronto.edu

**Prof. Eric Yu**  
eric@cs.toronto.edu

---

## Symbolic Optimization with SMT Solvers | BOOTH 42 (SE)



The rise in efficiency of Satisfiability Modulo Theories (SMT) solvers has created numerous uses for them in software verification, program synthesis, functional programming, refinement types, etc. In all of these applications, SMT solvers are used for generating satisfying assignments (e.g., a witness for a bug) or proving unsatisfiability/validity (e.g., proving that a subtyping relation holds). We are often interested in finding not just an arbitrary satisfying assignment, but one that optimizes (minimizes/maximizes) certain criteria.

SYMBA is an efficient SMT-based optimization algorithm for objective functions in the theory of linear real arithmetic (LRA). Given a formula  $\Phi$  and an objective function  $t$ , SYMBA finds a satisfying assignment of  $\Phi$  that maximizes the value of  $t$ . SYMBA utilizes efficient SMT solvers as black boxes. As a result, it is easy to implement and it directly benefits from future advances in SMT solvers. Moreover, SYMBA can optimize a set of objective functions, reusing information between them to speed up the analysis.


**Yi Li**  
liyi@cs.toronto.edu

**Prof. Marsha Chechik**  
chechik@cs.toronto.edu

## SOFTWARE ENGINEERING | ARTIFICIAL INTELLIGENCE (AI)

### A Mixed-Initiative Approach to Project Scheduling | **BOOTH 8 (SE + AI)**

Rishi Visual Lessor



A software tool for scheduling projects with multiple resource constraints was developed. A local search algorithm was applied to guide the user to optimal or near-optimal solutions. A mixed-initiative approach was used to allow for collaboration between the user and the software through a novel interactive visualization of the optimization problem.

**Kiran Sachdev**  
ksach@cs.toronto.edu

**Prof. Allan Borodin**  
bor@cs.toronto.edu

---

## SYSTEMS/NETWORKING (SysNet)

### A Simple Programming Model for Scalable SDN Applications | **BOOTH 46 (SysNet)**

Simplicity is a prominent advantage of Software-Defined Networking (SDN), and is often exemplified by implementing a complicated control logic as a simple control application deployed on a centralized controller. When it comes to practice, however, such simple control applications transform into complex logics on distributed control platforms, since they need to tolerate eventual consistency (as existing control platforms favor availability) and implement complex coordination and partitioning mechanisms (as existing control platforms envision that such functions belong to control applications, with the excuse of generality). As a result, distributed control applications are polluted with boilerplates of distributed programming that are usually more complicated than the control logic itself.


Here, we present a programming model that simplifies the development process of distributed applications. It is familiar and intuitive, yet generic enough to implement different communication patterns (such as Request/Response and Pub/Sub) and existing distributed controllers (such as ONIX and Kandoo). Moreover, we have implemented a highly efficient control platform and our evaluations indicate that the proposed programming model does not impose an inherent scalability bottleneck.

**Soheil Hassas Yeganeh**  
soheil@cs.toronto.edu

**Prof. Yashar Ganjali**  
yganjali@cs.toronto.edu

### Oasis: Energy Proportionality with Partial Server Consolidation | BOOTH 36 (SysNet)

Cloud data centers operate at very low utilization rates resulting in significant power waste. We introduce Partial Server Consolidation (PARSEC), a new approach to energy-oriented cluster management that achieves dense server consolidation with minimal network transfers.



**Junji Zhi**

zhij@cs.toronto.edu

**Prof. Eyal de Lara**

delara@cs.toronto.edu

**Nilton Bila**

IBM T.J. Watson Research Center

---

### Analyzing and Improving the Performance of Spatial Database Processing | BOOTH 1 (SysNet)



Spatial database workloads are very important because they are involved in a wide variety of spatial applications: from urban planning and land information management in general, to environmental applications such as flood risk impact, or toxic spill analysis, and even emerging new application domains like biomedical imaging and MRI data analysis. All of these applications use data with spatial attributes and perform heavy processing on data to obtain various reports and forecasts. Unfortunately, spatial query properties and spatial workload characteristics have not been well understood. We have identified

complex spatial queries which exhibit considerably poor performance in terms of execution time. This project focuses on analyzing spatial workloads from a performance perspective, identifying the performance bottlenecks, and proposing solutions for substantially speeding up spatial processing.

**Bogdan Simion**

bogdan@cs.toronto.edu

**Prof. Angela Demke Brown**

demke@cs.toronto.edu

---

### Traffic Statistics Collection with FleXam | BOOTH 17 (SysNet)

One of the limitations of wildcard rules in Software Defined Networks, such as OpenFlow, is losing visibility. FleXam is a flexible sampling extension for OpenFlow that allows the controller to define which packets should be sampled, what parts of each packet should be selected, and where they should be sent. Here, we present how FleXam enables the controller to dynamically adjust sampling rates and change the sampling scheme to optimally keep up with a sampling budget in the context of a traffic statistics collection application.

**Sajad Shirali-Shahreza**

shirali@cs.toronto.edu

**Prof. Yashar Ganjali**

yganjali@cs.toronto.edu

## THEORY

### **Linear Time Maximum Weighted Independent Set On Co-comparability Graphs | BOOTH 54 (Theory)**

The maximum weighted independent set (WMIS) problem is a widely studied problem as it arises naturally in different applications such as scheduling, combinatorial auctions, and molecular biology. The problem is NP-hard for arbitrary graphs; we restrict ourselves to the class of co-comparability graphs and present the first linear time algorithm to compute a WMIS directly on the given co-comparability graph. As a corollary, we get the minimum weight vertex cover of a co-comparability graph in linear time as well.

**Lalla Mouatadid**  
lalla@cs.toronto.edu

**Prof. Ekkehard Köhler**  
Brandenburgische Technische Universität  
ekoehler@math.tu-cottbus.de

---

### **Preference Elicitation, Approximate Stability, and Interview Minimization in Stable Matchings | BOOTH 38 (Theory)**

Algorithms for stable marriage and related matching problems typically assume that full preference information is available. 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. Furthermore, little work has investigated schemes for effectively eliciting agent preferences using either preference (e.g., comparison) queries or interviews (to form such comparisons); and no work has addressed how to combine both.

We describe the use of maximum regret to measure the (inverse) degree of stability of a matching with partial preferences; minimax regret to find matchings that are maximally stable given 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. We show that several of our schemes find stable matchings while eliciting considerably less preference information than Gale-Shapley.

We also develop a new model for representing and assessing agent preferences that accommodates both eliciting known preference information and (heuristically) minimizing the number of queries and interviews required to determine a stable matching. Our Refine-then-Interview (RtI) scheme uses coarse preference queries to refine knowledge of agent preferences and relies on interviews only to assess comparisons of relatively “close” options. Empirical results show that RtI compares favorably to a recent pure interview minimization algorithm, and that the number of interviews it requires is generally independent of the size of the market.

**Joanna Drummond**  
jdrummond@cs.toronto.edu

**Prof. Craig Boutilier**  
cebly@cs.toronto.edu