RESEARCH IN ACTION

Computer Science
UNIVERSITY OF TORONTO

MESSAGE FROM DEPARTMENT CHAIR

Research in Action is an opportunity to showcase to industry, prospective students, and the public, the exciting advances in computer science research. As a field, computer science is evolving, needed by every sector from banking, to entertainment, to law, to medicine. Computing will solve some of the world's greatest challenges, from helping cities cope with big data, to reducing our impact on the environment.

We have a long-standing history of collaboration with industry partners across Canada and abroad. Whether it's commercializing research developed here, receiving financial support for joint projects, or sponsoring a research internship through our Master of Science in Applied Computing program, the opportunities for our students and faculty to work together with industry to solve problems is limitless. If you are interested in learning more, then I encourage you to get in touch with our Associate Chair for Research & Industrial Relations, Khai Troung.

Research in Action also celebrates the work of our students, whether the project is in its earliest phase or near its end. Some inventions may become startups. Others will take their knowledge into the workplace, further graduate work, or teaching, all advancing U of T computer science's influence globally. The breadth and depth of these research examples are only the beginning.

Ravin Balakrishnan
Professor & Chair

MESSAGE FROM ASSOCIATE CHAIR, RESEARCH & INDUSTRIAL RELATIONS

I would like to take this opportunity to thank our talented students for presenting at Research in Action. Their work highlights the innovative research happening at the Department of Computer Science at the University of Toronto.

This year's Research in Action coincides with the department's Grad Visit Day – an event which gives prospective graduate students an opportunity to learn about our department, meet our faculty and current graduate students. We hope that these prospective graduate students and all of our guests have a chance to engage and learn from the projects presented at Research in Action.

I welcome, and am excited by, the possibility of exploring new collaborative and interdisciplinary efforts that will allow us to continue to make an impact and elevate our research and teaching. Please contact me if you have any interest in future collaborations.

Khai Truong
Associate Professor and Associate Chair, Research & Industrial Relations
acrir@cs.toronto.edu

A FRAMEWORK FOR MOBILE COLLABORATION BETWEEN SENIORS AND CAREGIVERS

Older Adults want and need increased opportunities to learn new things, grow socially and intellectually, and contribute their knowledge and talents to society. Participating in intellectually and socially complex activities has been shown to provide cognitive benefits to older adults. However, with the increasing complexity of knowledge, increasing size of caregiver circle, and the extensive reliance on computers, seniors demand to maintain their sense of control and independence in life to continue to learn, understand and grow socially and intellectually. The project is a theoretical framework for intelligent assistive technologies that offer the potential to fill this gap and provide practical support and autonomy to older adults while alleviating the burden on caregivers.

Yomna Aly
yomna.aly@mail.utoronto.ca

Professor **Cosmin Munteanu**
cosmin.munteanu@mail.utoronto.ca

APPLICATION OF NEURAL NETWORK TO EFFICIENT VALUATION OF LARGE VA PORTFOLIOS

Variable annuity (VA) products expose insurance companies to considerable risk because of the guarantees they provide to buyers of these products. Managing and hedging the risks associated with VA products requires intraday valuation of key risk metrics for these products. The complex structure of VA products and computational complexity of their accurate evaluation has compelled insurance companies to adopt Monte Carlo (MC) simulations to value their large portfolios of VAs. Because the MC simulations are computationally demanding, especially for intraday valuations, insurance companies need more efficient valuation techniques and the computational complexity of most of existing academic methodologies surpasses the computational requirements of MC simulations.

In this project, we present a spatial interpolation framework to significantly decrease the computational complexity of MC simulations. Our framework uses a neural network approach to find an effective distance function that is essential to the performance of spatial interpolation techniques. Unlike traditional spatial interpolation schemes that cannot provide all of the key

properties of accuracy, efficiency, and granularity, the proposed approach is accurate, efficient, and provides an accurate granular view of the input portfolio.

Amir Hejazi
amir@cs.toronto.edu

Professor **Kenneth Jackson**
krj@cs.toronto.edu

AUTOMATIC REAL-TIME MUSIC GENERATION

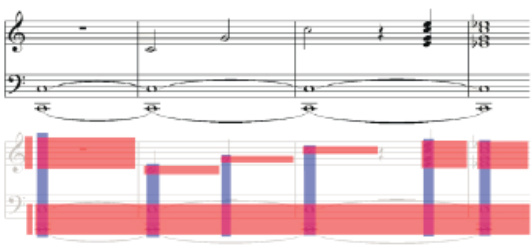

Music composition can be a challenge for many small- to medium-sized game companies, largely due to the expense and difficulty in creating original music tracks for each level of a game. To address this, we devised a tool that automatically generates original music for the different parts of a game, by training the music generator off source music whose style the game designer wishes to imitate. The generator then creates original music in that style in real-time, until the game signals to switch to another style of music. This software, first created in 2011, has been refined to produce music that is not just coherent, but also imitates the composer's larger music structure and allows for the incorporation of multiple pieces of music in the training.

Tiffany Tong
tiffany.tong@mail.utoronto.ca

Fabian Chan, Engineering Science
chanfabi@mail.utoronto.ca

Associate Professor, Teaching Stream, **Steve Engels**
sengels@cs.toronto.edu

BUDGETARY EFFECTS ON PRICING EQUILIBRIUM IN ONLINE MARKETS

We consider a pricing game with strategic vendors and a single buyer, modelling a scenario in which multiple competing vendors have very good knowledge of a buyer, as is common in online markets. In contrast to previous work we consider the realistic feature that the buyer has a fixed budget instead of unlimited funds.

When the buyer's valuation function is additive, we are able to completely characterize the different possible pure Nash Equilibria (PNE) and in particular obtain a necessary and sufficient condition for uniqueness. Furthermore, we characterize the market clearing (or Walrasian) equilibria for all submodular valuations. Finally, with respect to social welfare, while without budgets all equilibria are optimal (i.e., POA = POS = 1), we show that with budgets the worst equilibrium may only achieve 1/(n–2) of the best equilibrium.

Tyrone Strangway
tyrone@cs.toronto.edu

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

CORRELATION RISK

We investigate numerical and approximate solutions to option pricing problems with a non-constant correlation process.

Nat Leung
natleung@cs.utoronto.ca

Professor **Christina Christara**
ccc@cs.toronto.edu

DUALKEY: MINIATURE SCREEN TEXT ENTRY VIA FINGER IDENTIFICATION

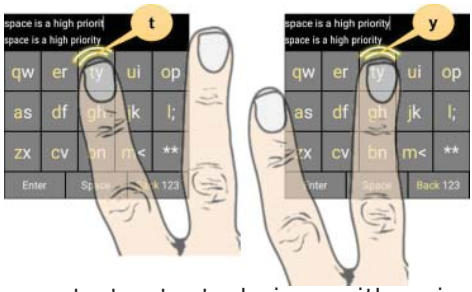

Fast and accurate access to keys for text entry remains an open question for miniature screens. Existing works typically use a cumbersome two-step selection process, first to zero-in on a particular zone and second to make the key selection. We introduce DualKey, a miniature screen text entry technique with a single selection step that relies on finger identification. We report on the results of a 10 day longitudinal study with 10 participants to evaluate speed, accuracy, and learning. DualKey outperformed the existing techniques on long-term performance with a speed of 19.6 WPM. We then optimized the keyboard layout for reducing finger switching time based on the study data. A second 10 day study with eight participants showed that the new sweqty layout improved upon DualKey even further to 21.59 WPM for long-term speed, was comparable to existing techniques on novice speed and outperformed existing techniques on novice accuracy rate.

Aakar Gupta
aakar@cs.toronto.edu

Professor **Ravin Balakrishnan**
ravin@dgp.toronto.edu

ENERGY EFFICIENT ILLUMINATION AND IMAGING

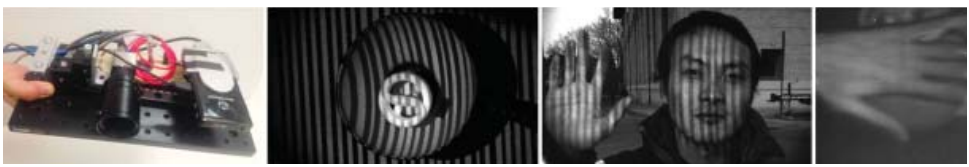

We present our energy-efficient structured light camera: a device for projecting a known light pattern onto environments containing strong ambient light sources and indirect transport effects. Structured light devices use patterns with spatial structures (Kinect and Leap Motion sensors) and temporal structures (time-of-flight depth sensors) towards applications such as 3D scanning. However, conventional sensors tend to be restricted to indoor environments, given that strong ambient light sources like the

sun completely overwhelm the low-power light signal emitted by such devices. Moreover, these devices often assume that light emitted into a scene reflects directly back to the camera; any indirect light paths (caustics, diffuse interreflections) can severely confuse such devices. Our camera demonstrates near-optimal energy-efficient imaging for 3D scanning through a novel combination of a low-power laser projector and a rolling shutter camera. We demonstrate our camera's ability to (1) capture structured light images of very bright scenes – even a light bulb that has been turned on; (2) capture direct (and18 indirect-only) video with optimal energy efficiency; and (3) record structured light images without interference from other devices.

Mian Wei
mian.wei@mail.utoronto.ca

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

GROWINGLEAF: MODELING AND ANALYSIS FOR GOALS WITH TEMPORAL DYNAMICS

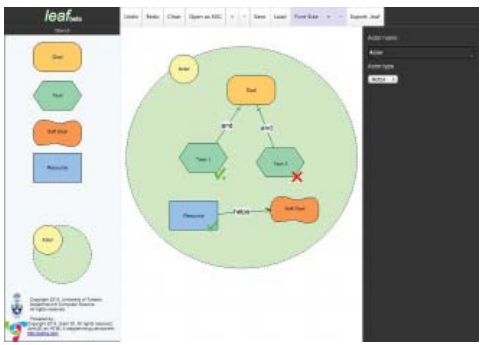

GrowingLeaf is a web-based tool for modeling and analyzing goal models with temporal dynamics. Goal models for early phase requirements enable modelers to elicit stakeholders' intentions, analyze dependencies, and select preferred alternatives. Standard analysis techniques provide options for analysis of static goal models but do not consider the dynamic environment that the model represents and do not evaluate the intentions over time. Using symbolic and concrete simulation techniques, we provide tooling that enable stakeholders to choose between design alternatives, ask what-if questions, and plan for software evolution in an ever-changing world. If you have evolving requirements, see what insight our tool can provide.

Alicia M. Grubb
amgrubb@cs.toronto.edu

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

IMPROVING POWER FOR DETECTING DISEASE GENES USING BIOLOGICALLY INFORMED GRAPHICAL MODELS

Identifying genes associated with complex human diseases is one of the main challenges of human genetics and computational medicine. In the context of rare variants, there are several approaches that aggregate the signal of all variants in a gene to statistically test its association with a disease. However, these approaches usually consist of univariate tests where the complementarity between genes is not taken into account. To increase the power of identifying genes associated with diseases, we developed a method based on a hierarchical graphical model jointly incorporating all genes and all variants. Our method integrates external biological information such as variant harmfulness predictions and gene networks as priors over the model's variables.

Our extensive simulations show that our method has up to one order of magnitude more power to identify disease genes compared to existing variant aggregating methods. When run on several exome sequencing datasets, our approach was able to identify genes previously associated with lipid traits with a fraction of samples compared to existing methods. Finally, our method produces statistics at the level of individual patients, such as which genes and variants are important in that individual case.

Aziz M. Mezlini
aziz.mezlini@utoronto.ca

Professor **Anna Goldeberg**
anna.goldenberg@utoronto.ca

LEARNING LEXICAL EMBEDDINGS WITH SYNTACTIC AND LEXICOGRAPHIC KNOWLEDGE

We propose two improvements on lexical association used in embedding learning: factorizing individual dependency relations and using lexicographic knowledge from monolingual dictionaries. Both proposals provide low-entropy lexical co-occurrence information, and are empirically shown to improve embedding learning by performing notably better than several popular embedding models in similarity tasks.

Tong Wang
tong@cs.toronto.edu

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

OPPORTUNISTIC STORAGE MAINTENANCE

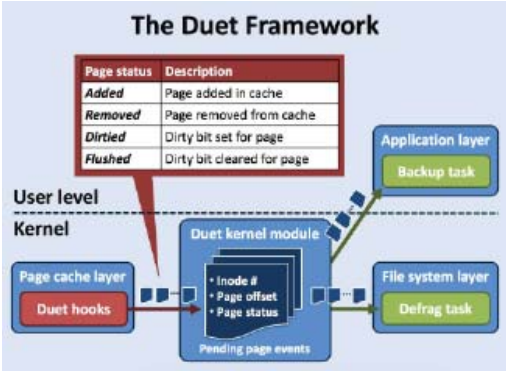

Storage systems rely on maintenance tasks, such as backup and layout optimization, to ensure data availability and good performance. These tasks access large amounts of data and can significantly impact foreground applications. We argue that storage maintenance can be performed more efficiently by prioritizing processing of data that is currently cached in memory. Data can be cached either due to other maintenance tasks requesting it previously, or due to overlapping foreground I/O activity.

We present Duet, a framework that provides notifications about page-level events to maintenance tasks, such as a page being added or modified in memory. Tasks use these events as hints to opportunistically process cached data. We show that tasks using Duet can complete maintenance work more efficiently because they perform fewer I/O operations. The I/O reduction depends on the amount of data overlap with other maintenance tasks and foreground applications. Consequently, Duet's efficiency increases with additional tasks because opportunities for synergy appear more often.

George Amvrosiadis
gamvrosi@cs.toronto.edu

Professor **Angela Demke Brown**
demke@cs.toronto.edu

Professor **Ashvin Goel**, Electrical and Computer Engineering
ashvin@eecg.toronto.edu

PHENOTYPE COMPARISON VISUALIZATIONS OF PATIENT COHORTS

In medicine, phenotypes are the traits of a patient that describe the signs, symptoms, and manifestations that characterize a given disease. Certain combinations of phenotypes are more commonly observed in specific rare genetic disorders. Genetics researchers investigate the variation of phenotypes between diagnosed patients of rare genetic diseases to broaden our understanding of genetic variations. We are developing a visual analytics tool that supports genetics researchers in this task. By leveraging a top-down approach to comparing phenotypes within patient cohorts we enable both the exploration of phenotypic variation between cohorts of patients and between patients within each cohort.

Michael Glueck
mglueck@dgp.toronto.edu

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

POROUS FLUID SIMULATION WITH HETEROGENEOUS ELEMENTS

Most solid matter is porous, but interactions with fluids in computer graphics often overlook this important property. We present a FLIP-based method for simulating fluid interactions with porous objects with little extra cost by using Discrete Exterior Calculus. We encode both object boundaries and permeability with a Hodge dual to produce a porosity-aware Laplacian operator.

Michael Tao
mtao@dgp.toronto.edu

Professor **Eugene Fiume**
elf@dgp.toronto.edu

SAT IS AN EFFECTIVE AND COMPLETE METHOD FOR SOLVING STABLE MATCHING PROBLEMS WITH COUPLES

Stable matchings can be computed by deferred acceptance (DA) algorithms. However, such algorithms become incomplete when complementarities exist among the agent preferences: they can fail to find a stable matching even when one exists. We examine stable matching problems arising from labour market with couples (SMP-C). The classical problem of matching residents into hospital programs is an example. Couples introduce complementarities under which DA algorithms become incomplete. In fact, SMP-C is NP-complete.

Inspired by advances in SAT and integer programming (IP) solvers, we investigate encoding SMP-C into SAT and IP and then using state-of-the-art SAT and IP solvers to solve it. We also implemented two previous DA algorithms. After comparing the performance of these different solution methods, we find that encoding to SAT can be surprisingly effective, but that our encoding to IP does not scale as well. Using our SAT encoding, we are able to determine that the DA algorithms fail on a non-trivial number of cases where a stable matching exists. The SAT and IP encodings also have the property that they can verify that no stable matching exists, something that the DA algorithms cannot do.

Joanna Drummond
jdrummond@cs.toronto.edu

Andrew Perrault
perrault@cs.toronto.edu

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

Professor **Fahiem Bacchus**
fbacchus@cs.toronto.edu

SECONDSKIN

SecondSkin is a sketch-based modelling focused on the creation of layered structures, comprised of shape interdependent 3D volumes. Our approach is built on three novel insights gleaned from an analysis of sketches. First, we observe that a closed loop of strokes typically define surface patches that bound volumes in conjunction with underlying surfaces. Second, a significant majority of these strokes map to a small set of curve-types, that describe the 3D geometric relationship between the stroke and underlying layer geometry. Third, we found that a few simple geometric features allow us to consistently classify 2D strokes to our proposed set of 3D curve-types. With these insights, our system infers 3D curves from 2D sketched strokes and automatically surfaces closed loops allowing users to construct layered geometry.

Christopher De Paoli
chrisdepaoli@gmail.com

Karan Singh
karan@dgp.toronto.edu

SIMPLE SEARCH ALGORITHMS ON SEMANTIC NETWORKS LEARNED FROM LANGUAGE USE

Recent empirical and modeling research has focused on the semantic fluency task because it is informative about semantic memory. An interesting interplay arises between the richness of representations in semantic memory and the complexity of algorithms required to process it. It has remained an open question whether representations of words and their relations learned from language use can enable a simple search algorithm to mimic the observed behavior in the fluency task. Here we show that it is plausible to learn rich representations from naturalistic data for which a very simple search algorithm (a random walk) can replicate the human patterns. We suggest that explicitly

structuring knowledge about words into a semantic network plays a crucial role in modeling human behavior in memory search and retrieval; moreover, this is the case across a range of semantic information sources.

Filip Miscevic
filip.miscevic@mail.utoronto.ca

Aida Nematzadeh
aida@cs.toronto.edu

Professor **Suzanne Stevenson**
suzanne@cs.toronto.edu

SKETCHSOUP: EXPLORATORY IDEATION USING DESIGN SKETCHES

A number of quick-and-dirty sketches that capture design inspirations, model variations, and alternate viewpoints of a visual concept are a hallmark of early-stage conceptual design. We present SketchSoup, a workflow that allows designers to explore the design space induced by such sketches. We take an unstructured collection of sketch images as input, register them using a multi-image matching algorithm, and present them as a 2D interpolation space, suitable for interactive design exploration. Such "sketch tourism" allows designers and their patrons to make better informed choices, and can provide inspiration for further drawings, that seamlessly feed back into the design space as additional image inputs. Our contribution is thus a solution that fills a significant gap in the early ideation stage of conceptual design, by judiciously combining and adapting various image processing techniques to the sketch domain, where the images are dominated, not by color or texture, but messy and imperfect stroke contours.

Rahul Arora
arorar@cs.toronto.edu

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

SOQR: SONICALLY QUANTIFYING CONTENT LEVELS INSIDE A CONTAINER

Sensing the amount of content in a container is an important ubiquitous computing research challenge with a wide variety of practical applications. It can provide an understanding of when an item has been used, how much content has been consumed and whether a product needs to be replenished. Such knowledge, for instance, can help patients manage their medication compliance or remind care-givers when they may need to refill patients' medication bottles. Many existing approaches may be impractical or expensive. To tackle this problem, we have designed a small acoustic sensor that can be attached to an external wall of a container and sense the content level inside it.

Mingming Fan
mfan@cs.toronto.edu

Professor **Khai Truong**
khai@cs.toronto.edu

SYNTHESIZING FINITE AND INFINITE PLANS WITH TEMPORALLY EXTENDED GOALS

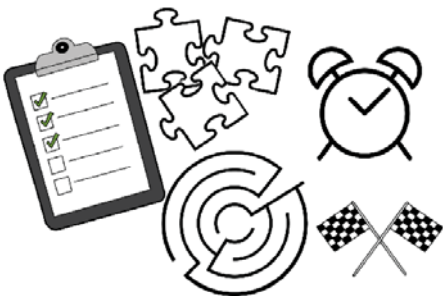

Automated planning is a branch of AI that concerns the generation of a set of ordered actions, or a policy, for execution by some agent or agents in order to achieve a stated objective. Automated planning has applications in problems as diverse as logistics planning, robot planning, software verification and testing, and narrative generation. Our concern is with problems where actions can be non-deterministic, and where plans can be finite or infinite. Our focus is on plan objectives involving properties of the plan trajectory, rather than simply a goal state, including maintenance or realization of various properties or the mandated execution of particular actions in a particular order. Such temporally extended goals can be expressed in Linear Temporal Logic (LTL). To realize non-deterministic planning with LTL goals, our approach translates LTL goals into either alternating or non-deterministic finite state automata, and exploits a state-of-the-art fully observable non-deterministic planner to compute solutions. Our system addresses a diverse spectrum of LTL planning problems that to this point had not been solvable using AI planning techniques. We do so while demonstrating competitive performance relative to the state of the art

in LTL planning.

Alberto Camacho
acamacho@cs.toronto.edu

Christian Muise
Massachusetts Institute of Technology

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

Eleni Triantafilou
eleni@cs.toronto.edu

Jorge Baier
Pontificia Universidad
Católica de Chile

TEXT AND SPEECH PROCESSING FOR THE DETECTION OF DEMENTIA

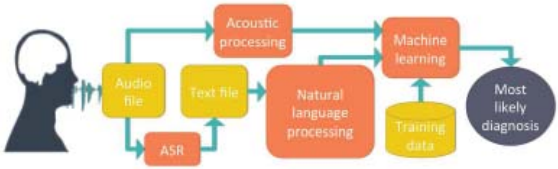

Dementia is a gradual decline of cognitive abilities, often resulting from neurodegeneration. In some cases, such as primary progressive aphasia (PPA), language abilities are specifically impaired. In other cases, such as Alzheimer's disease, language decline may occur together with other cognitive impairments. In each of these instances, a narrative language sample can provide a wealth of data regarding an individual's linguistic capabilities. Traditionally, analysis of speech samples was conducted by hand, but this is painstaking and time-consuming work. We extract hundreds of relevant lexical, syntactic, and acoustic features from short speech samples, and use machine learning classifiers to distinguish between speakers with dementia and older healthy controls with over 80% accuracy. This technology has numerous potential applications, including screening, quantifying the effect of interventions, longitudinal assessment over time, and early detection.

Kathleen Fraser
kfraser@cs.toronto.edu

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

Professor **Frank Rudzicz**
frank@cs.toronto.edu

Jed Meltzer
Rotman Research Institute

Elizabeth Rochon
Department of Speech-Language Pathology

THE ALLT PROJECT (ACCESSIBLE, LARGE-PRINT, LISTENING AND TALKING E-BOOK) AND DIGITAL STORYTELLING

When looking at the social aspects of reading or storytelling, one can observe that accessibility is part of reading together; for example, an adult reading to their older parent that suffers from age-related impaired eyesight. The ALLT project (an Accessible, Large-print, Listening and Talking e-book reader) goes beyond enabling access to a read-aloud book, and allows a collaborative reading experience in which people read together, with the recorded audio synchronized with the book content and later played back (in a familiar voice) by the visually-impaired user.

Our research is now branching off this project to explore collaborative digital storytelling between generations, but using family photographs instead of an e-book. This tool will create persistent multimedia stories as told by seniors and will be shareable, with summary captions for each photo derived from the oral stories. We hope that this tool can work as a trigger for storytelling by being enjoyable to use and by having shareable outputs (i.e., multimedia stories and photo captions). We expect this solution to increase socialization among seniors and preserve family knowledge without undue effort.

Rachel Benett Axtell
axtell@mail.utoronto.ca

Professor **Cosmin Munteanu**
cosmin@taglab.ca

UNDERSTANDING ASSISTIVE TECHNOLOGY'S EFFECT ON RELATIONSHIP MAINTENANCE

Assistive technologies are devices designed to help their users interact with the world at large by replacing or supplementing lost functions often brought on by a progressive illness or a disability. In this project, we investigate the effectiveness of assistive technologies in helping their users interact and maintain relationships specifically with people in their lives. In particular, progressive illnesses and disabilities may sometimes thrust a patient's family members or friends into the role of informal caregiver. This new role can potentially change the relationship between the patient and their informal caregiver. Thus, we are interested in exploring the how assistive devices are used as a relationship maintenance tool in patient/caregiver relationships as well as the impact of these devices on the change in relationship.

Jinglan Lin
cjlin@cs.toronto.edu

Karyn Moffatt
kmoffatt@cs.toronto.edu

Professor **Khai Truong**
khai@cs.toronto.edu

WEARABLE SENSOR-BASED FALL RISK ASSESSMENT AND ACTIVITY TRACKING FOR OLDER ADULTS: A STUDY OF EFFECT

Seniors are faced with the risk of falling. Identifying such risk is crucial for implementing appropriate intervention to prevent future falls. In this project, a custom wearable sensor device is used to identify the risk of falling and provide activity data for the elderly user and caregivers. This device will be field tested to reveal its effectiveness, acceptance, and whether seniors and caregivers regard such implement helpful in preventing falls.

Yusheng Wu
yusheng.wu@mail.utoronto.ca

Professor **Ronald Baecker**
ron@taglab.ca

WEARCOPD: MONITORING COPD PATIENTS REMOTELY USING SMARTWATCHES

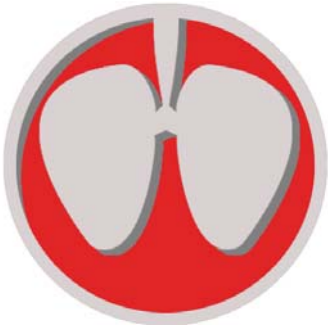

Chronic Obstructive Pulmonary Disease (COPD) is a chronic lung disease that is characterized by airway obstruction, coughing, shortness of breath and sputum production. In 2009, the Public Health Agency of Canada reported that 772,200 Canadians suffer from COPD. The World Health Organization estimated that over 3 million people died of COPD in 2012 and ranks it as the 4th leading cause of death, tied with HIV/AIDS.

An exacerbation of COPD is a sudden worsening of the disease. Exacerbations result in more frequent and severe coughing and more difficulty breathing. If an exacerbation is not treated quickly, hospitalization may be required which is expensive and decreases the patients quality of life.

WearCOPD is a system we are developing that can remotely monitor patients. Patients are equipped with an Android Wear smartwatch and Android smartphone and are free to go about their daily activities. WearCOPD collects audio traces, heart rate data, accelerometer data, a step count and questionnaire responses. Using this data, we hope to monitor changes in patients conditions over time and provide early warning of an oncoming exacerbation.

Daniyal Liaqat
dliaqat@cs.toronto.edu

Professsor **Eyal de Lara**
delara@cs.toronto.edu

DEPARTMENT OF COMPUTER SCIENCE
University of Toronto
Sandford Fleming Building
10 King's College Road, Room 3302
Toronto, Ontario M5S 3G4
www.cs.toronto.edu