



ARIA

Applied Research in Action

Master of Science in Applied Computing
(MScAC) Program

November 27, 2019



Message from the Chair

Today, all aspects of our lives are impacted by technological innovations, so much so that technology weaves through nearly every action we take. It is little wonder that the demand for highly-skilled and highly-trained computational scientists continues to increase rapidly.

Toronto has emerged as a world leader in incubating and growing new technology firms. The Department of Computer Science at the University of Toronto is committed to being an integral player in this phenomenon, ensuring that our teaching and learning environment leverages the many technological developments occurring in broader society, while building bridges to industry to ensure our research has the greatest possible impact.

The Master of Science in Applied Computing (MScAC) program has emerged as a critical component of our outreach effort. The program is premised on a partnership model that builds industry linkages to ensure these students get the best of both academic and industry research training. We are proud of the contributions these students are already making in tackling local and global challenges, as we see them contributing to building healthy, sustainable communities.

The MScAC program is part of our broader efforts to create engagement opportunities with industry partners across all of our graduate and research programs.

This year, many of our industrial participants are joining the showcase with booths displaying their work, and providing opportunities for faculty and students across the University of Toronto and, particularly in the Department of Computer Science, to engage in discussions about research partnerships.

Finally, the Department of Mathematics, the Department of Statistical Sciences and the Department of Physics have actively joined the MScAC program and we welcome them to ARIA as well.

For more information on research partnerships or to learn more about our research programs, please contact us at acrir@cs.toronto.edu.

I wish all of our guests at ARIA a stimulating and enjoyable time learning about this year's internship projects of our MScAC students and about exciting industrial opportunities.

Marsha Chechik

*Professor & Chair
Department of Computer Science*



Message from the Program Director

Welcome to ARIA 2019, our Applied Research in Action showcase, which highlights the exciting innovations being developed by the students in our Master of Science in Applied Computing (MScAC) program.

We are delighted that you are joining us in celebrating the remarkable contributions of these incredible students, to developing world-class solutions to the challenges posed by our industry partners.

For the past 15 months, these MScAC students have been immersed in a unique learning environment that fuses academic training with industry research. The program is designed to give them advanced technical skills and to create a platform for applying these skills in creative ways. Today we are seeing first-hand the end-result of this process. ARIA is a milestone event in their degree progression - a time to honour their hard work, commitment and dedication. We also want to recognize the extraordinary group of companies who have contributed so much to our students' career successes.

Our program continues to evolve and we are pleased to know be offering two specialized concentrations with MScAC. The Data Science (DS) concentration (joint between the Department of Computer Science and the Department of Statistical Sciences) focuses on extracting and inferencing information from complex data sets. The Applied Math concentration (joint between the Department of Computer Science and the Department of Mathematics) focuses on the application of advanced quantitative and mathematical modeling techniques in a real-world setting.

This year a number of our industry partners are displaying their own R&D efforts. A thank you to Autodesk Research, BlueCat Networks, Deloitte Canada, Geotab, Layer 6 AI, SOTI Inc., Vanguard

Investments and Velocity for their wonderful displays. And a special thanks to Mitacs for their generous support to the MScAC program since its inception.

In February 2020 we will hold our Applied Research Internship Expo (ARIE), where our industry partners meet and start a collaborative rapport with our 2019/20 cohort. If you are interested in hosting a research intern next year, please contact us at miscac@cs.toronto.edu.

Finally, we thank the many people who continue to make the MScAC program and ARIA such a success. First and foremost is Claire Mosses, who so adeptly ensures all the pieces are in place for student success. Ryan Perez did much of the heavy lifting in putting together ARIA while Dina Savana provided significant logistical support for this event. A special thanks to the entire team in the Department of Computer Science for their continued contributions to the MScAC program and ARIA. Once again, congratulations to our 2018/19 Applied Computing cohort!

Arvind Gupta

*Professor
Academic Director, Professional Programs
Department of Computer Science*



About the Keynote Speaker

Amber Mac

Speaker, Host, Author, Entrepreneur

Amber Mac(Arthur) started her career in San Francisco and Boston during the dot-com boom. As a strategist for Razorfish and Director of Marketing for an e-procurement software company, she spent four years in the technology start-up trenches. Amber left the start-up world to join Microsoft to build one of the first female-focused lifestyle portals. In 2006, she started her own digital agency. The company's first client was world-renowned business coach, Tony Robbins, and now includes Microsoft, Google, GE, Paypal, Nintendo, Canada Goose, and Fast Company, and many more leading organizations. In 2010, she wrote the national bestselling business book Power Friending (Portfolio: New York). In 2016, she co-wrote Amazon bestseller, Outsmarting Your Kids Online. Most recently, she started co-hosting a new podcast series, presented by Accenture, called Marketing Disrupted.

Amber has keynoted more than 300 events around the world and has moderated sessions with: Canadian Prime Minister Justin Trudeau; former FBI Director James Comey; business coach Tony Robbins; famed astronaut Chris Hadfield; and many other notable leaders.

She is a regular business host and expert for Fast Company, CNN, Bloomberg, CBS, BNN, CTV, The Marilyn Denis Show, and SiriusXM Radio. In fact, one of the MScAC's first internships, My Voice, featured on her show, App Central. Plus, in 2018, she was named one of DMZ's 30 inspirational women making a difference in tech. In 2019, her podcast The AI Effect won Best Technology Series.

She is now recognized internationally as an innovation and technology leader.



2019 Industry Exhibitors



Deloitte.

layer6

GEOTAB[®]
management by measurement

SOTI[®]

Vanguard[®]


VELOCITY

Special Thanks to our 2019 Industry Project Partners

*Ada
Bridge 7
Conflux
CPPIB
DB Schenker
Delphia
Ecobee
Epson
Ethoca
Insitro
Kindred
Layer 6 AI
Li Ka Shing Centre
Loblaw Companies
Loblaw Digital
ModiFace
Neuroradiology Dept., Toronto Western
ODAIA*

*Pelmorex: Data Team
Pelmorex: The Weather Network
Phenomic AI
Properly
Quantum Black
RockMass Technologies
RSVP AI
Samsung AI
Scotiabank
SOTI
SSC & DST
Surgical Safety Technologies
Tactical Labs
Uber ATG
Ukko Agro
Untether AI
Valsys*

SSL Offloading Or Aharoni

SOTI MobiControl is an Enterprise Mobility Management (EMM) solution for all mobile devices and IoT endpoints. It is a mobile security and device management that delivers enhanced EMM features like; application management, content management, location services and remote help over wired or wireless networks. The solution ensures that all data on the device, or moving between the device and the company servers, remains private and secure.

The devices communicate with MobiControl through cryptographic protocol called Secure Socket Layer (SSL). Currently SOTI's server handles all SSL related load, as well as management server workload, on the same servers. Both the SSL processing, and the processing on management operations, are performed inside the same physical process. This results in increased memory

and CPU consumption.

The goal of the project is to identify performance issues with the current package distribution and file transfer architecture, and to make architectural improvements that result in decreasing server CPU utilization and more optimal network bandwidth, by isolating SSL computation into a separate process and offloading processing to another SSL device. The project breaks into several phases; understanding MobiControl system, building a testing environment that includes metrics and logs, building "SSL re-encryption" that gets an encrypted message that comes from the device, decrypting the message and encrypting it again, sending it to the server, and "SSL Offloading," which removes the encryption and decryption from the server and finds a solution to pass the certificate to the management server.

*Academic Supervisor (AS): David Lie
Industry Supervisor (IS): Lorne Wilkinson*

End-to-End Video Object Segmentation Module Abdalla Ahmed

Semi-supervised video object segmentation is the task of predicting a spatiotemporal segmentation volume for a video clip, conditioned on one single-frame reference annotation for each object track in the sequence. We take advantage of the prior knowledge from image-level datasets like COCO, whilst also using VOS datasets, like YouTube-VOS, to learn end-to-end merging of per-frame predictions. In doing so, we

demonstrate the importance of video object segmentation in using prior knowledge learned from more easily obtainable static datasets. We also propose a merging algorithm, which gathers object proposals and their corresponding semantically segmented masks, and merges them to produce accurate pixel masks for the video sequence.

*AS: Arvind Gupta & Sanja Fidler
IS: Irina Kezele*

SOTI Blockly: Block-based Programming for Enterprise Mobility Management Jamie Beverley & Sahil Narula

Block-based programming languages and other structural editors have appreciated popularity over text-based languages for novice programmers by eliminating syntax errors and other technical barriers. Existing block-based programming frameworks, such as Blockly and Scratch, have been designed with younger user demographics in mind, often for use in educational contexts. However, adult users (as well as experienced programmers) could benefit from structural programming editors that reduce errors by only permitting the construction of syntactically valid code structures.

SOTI Snap is a mobile application development tool that enables users without any experience in software development to build custom applications. However, where custom scripting logic is needed within a Snap application, the user is expected to know JavaScript. This research seeks to address the technical barriers faced by Snap users through the design and

implementation of a block-based scripting environment suited specifically to the needs of business professionals who do not have prior programming experience.

Four focus groups (N=40) were conducted over two months, in which participants were asked to use and comment on different prototypes of a block-base scripting environment embedded into SOTI Snap. Participants consistently emphasized the importance of familiar natural language terminology, effective onboarding and documentation, and a desire to test scripts quickly within the editor. The research culminated in the implementation of SOTI Blockly, a block-based programming environment for SOTI Snap. A usability study was conducted (N=12) to evaluate the design, motivating future work to explore a block search feature, context-aware block suggestions, and drag-and-drop usability improvements.

*AS: Ishtiaque Ahmed & Khai Truong
IS: Shash Anand*

Fast and Robust Object Detection with Part-Based Keypoint Priors

Vineeth Bhaskara

The Toronto AI Centre (AIC) is a key part of Samsung's global AI research efforts that aim to conduct long-term research, with potential for direct impact on Samsung products reaching hundreds of millions of users worldwide. The focus of the lab spans a broad spectrum of research areas, including computer vision, computational linguistics, mobile, and distributed intelligence. A key application of computer vision is object detection, which has many practical use cases. Detection systems have improved significantly over the years, especially with the help of deep learning.

Our work builds on a recent single-stage architecture that summarizes objects as keypoint detections of bounding box centres, and regresses the dimensions

AS: Roger Grosse

IS: Alex Levinshtein & Kosta Derpanis

(height and width) in a fully-convolutional way. This method is highly sensitive to errors in localizing centre keypoints. We propose auxiliary part-based tasks that allow the network to accumulate evidence from multiple parts (e.g., bounding box corners) for robust object localization. Specifically, we use multi-task learning with an additional training objective to estimate corners as keypoints and their associations as vector fields, where each pixel votes for the relative directions of keypoints. We additively fuse the part-level hidden features and predictions into the primary detection branch for better contextual information in localizing the centre. Experiments show an improvement of up to 10% (~ 3 mAP) over the baseline on MS COCO with a minimal runtime overhead.

Filtering Human-Chatbot Conversations by Sent

Paul Briggs

Ada Inc is one of the fastest growing startups in the chatbot space. They differentiate themselves from other companies in this competitive market by working directly with customer service teams. They provide their clients with the tools and support to design and deploy their own chatbot. Ada's clients can automate up to more than 70% of incoming customer queries and escalate users directly to live agents if the bot fails to address the customer's needs. Ada provides their clients with access to all customer-chatbot conversations. Interviews with clients and user behavior patterns identified that clients will spend 4-5 hours a week auditing bot conversations, with the intent of finding ways to improve their bot.

Furthermore, when doing so they're seeking "bad conversations" because they are often excellent sources of bot-improvement insights. However, conversations are sorted by date and, as such, the majority of conversations they read in that 4-5 hour period do not yield any useful insights. To fix this, this project proposes using a three stage, SVM-based, sentence-level sentiment analysis model to identify and surface "bad conversations." In doing so, we can greatly increase the efficiency with which our clients can make improvements to their bot and reduce the amount of tedious work required.

AS: Graeme Hirst

IS: Gordon Gibson

Average Reward Reinforcement Learning for System Optimization in Robotics Application

Bryan Chan

Kindred's SORT is a piece-picking robot that grasps, scans, and stows items using various AI and control techniques. It leverages learning techniques to adapt to the constantly changing environment of a retail fulfillment center. The scan phase is a critical step in the sortation loop, as the system attempts to identify the grasped item through scanning a barcode located on the item. A successful scan provides item information and its stow location, whereas a scan failure results in restarting the sortation loop, thus inducing significant block time. As a result, one major Key Performance Indicator (KPI) for SORT is scan throughput, which measures the number of scanned units per hour (UPH). This research project focuses on improving scan throughput.

AS: Florian Shkurti & Animesh Garg

IS: James Bergstra

One approach to improving scan throughput is to identify bad grasps and unscannable items. Historically, scan failures are due to barcodes being obscured, bent, or wrinkled. Scan throughput can be improved by identifying these scenarios and restarting the sortation loop earlier. The proposed method aims to optimize the overall sortation loop using Average Reward Reinforcement Learning, a learning technique for optimizing continual sequential decision-making problems. It dynamically learns the time trade-off between scanning and regrasping as it terminates the scan phase early. The two resultant models improved the scan throughput from 470 UPH to 500 UPH on one machine, and 570 UPH to 600 UPH on the other.

Tizen OS Support for SOTI MobiControl Interoperability

Yin-Hung Chen & Qi Zhao

Enterprises are seeking to improve business performance, by monitoring the health and safety of employees and increasing productivity, with the introduction of IoT devices. However, the lack of security and management are barriers affecting the adoption of these devices. SOTI MobiControl is an Enterprise Mobility Management (EMM) solution that secures and manages devices on different platforms, throughout their entire lifecycle. Our internship project is to expand SOTI MobiControl, by adding support for the Tizen Operating System. This is an operating system that is widely used by Samsung IoT devices because of its lighter weight and faster speed.

In this project, we tried to research the most effective approach to port the existing MobiControl agent from other platforms to

AS: Eyal de Lara

IS: Lorne Wilkinson

the Tizen platform. First, we investigated Tizen using its official documentation to understand the novel system and security design. Second, a survey of the existing SOTI products helped us understand the architecture of existing agents on various platforms. We determined port based on the Linux agent at last, and implemented necessary modifications on it.

Challenges we have faced so far include: an operating system with strict security protection, a light weight system with strict memory, CPU and battery limitations, and different versions of hardware and firmware. Our contribution has made the porting successful. The Tizen MobiControl agent is now able to run on different models of Samsung IoT devices and perform most of the functionalities.

Forecasting Company Fundamental Factors

Tianxiao Chen & Ranjani Murali

A more general framework for financial modeling is provided by machine learning, as opposed to its linear parametric predecessors, generalizing archetypal modeling approaches such as factor modeling and derivative pricing. As shown in previous working papers, the returns in stocks can be effectively captured via prediction using the fundamental-type information, instead of direct prediction of returns. The factors chosen have a

strong correlation with the returns. This rationale is strengthened by the logic that the patterns captured, by forecasting future fundamentals that are likely to be persistent even out-of-sample than in returns. This research work is designed to create a forecasting model through a more representative and diverse dataset exploration, to predict the company's future fundamentals using the LightGBM model.

AS: Anthony Bonner

IS: Jonathan Briggs

ICR, OCR and Machine Learning Based Corrections of Data Transcription from Scanned Business Documents

Chen Chen

SS&C is a financial technology company that specializes in handling the back office (e.g. regulatory compliance, accounting, etc.) of investment firms and banks. As SS&C is the US market leader in its field, it is unique in the quantity and variety of business documents it sees. The processing of these documents requires large amounts of manual labor and the automation of this process would result in large savings, due to the scale of SS&C's operations. This project focused on the image enhancement and character recognition components of this process. Efforts in the past to automate text recognition have consistently been hindered by factors such as the variability and diversity of text, complexity and interference of backgrounds, and imperfect imaging conditions. In addition, it is difficult to leverage SS&C's large amounts of data for deep learning approaches, due to

privacy regulations and a lack of key bits of metadata.

Work on this project decomposed the challenge of text recognition into three components: (1) create a data pipeline to synthesize a training dataset, which would be used to train (2) text localization models, and (3) text recognition models. For (1), an exhaustive search was conducted of all publicly available text datasets, which resulted in the creation of a five million field dataset that served as both an OCR and text localization dataset. For (2), fully convolutional neural networks were used to draw bounding boxes around text of interest, as well as provide a pixel level mask to isolate text. Finally, for (3), a combination of convolutional neural networks, recurrent neural networks, and attention was used to convert preprocessed images into text data.

AS: Joseph Jay Williams

IS: Paul Ives

Applying Machine Learning to Predict Demand Transference

Zi Yi Chen

Loblaw Companies Limited is Canada's food and pharmacy leader, with a network of corporate and independently-operated stores in communities across the country. Loblaw's purpose – Live Life Well – supports the needs and well-being of Canadians, who make one billion visits each year to the company's stores. Loblaw operates five independent divisions: Market, Shoppers Drug Mart, Discount, PC Financial and Joe Fresh.

The objective of this project is to research, design, build and experiment with a machine learning (ML) model that can determine where demand is transferred when a product isn't available to

be purchased/out of stock (OOS). Ultimately, in an OOS situation, the ML model for demand transference allows us to increase the stock of substitute product, ensuring low product waste and high customer satisfaction. The factors that determine demand transference may include the brand and price of the products, type and location of the stores, inventory level, consumer loyalty, calendar, competition, item features and trends. The demand transference ML model can be applied to various businesses across the organization to improve our sales forecast, inventory policy, online order fulfillment, and assortment.

AS: Qiang Sun

IS: Miti Modi

Machine Learning Aided Detection of Brain Aneurysms

Daniel Dastoor

Brain aneurysms are weaknesses found in the walls of cerebral blood vessels, generating outpouchings or sacs. These are relatively common, occurring in 2-5% of the general population. If they rupture, they can cause a severe condition where up to one-third of cases may be fatalities, one third may result in severe debilitation, and one third may result in recovery through urgent treatment. As such, the detection of brain aneurysms, for example from scans such as computed tomography angiograms (CTAs), is important. However, these occurrences can be difficult to find on scans – especially when small or close to the bone. Nonetheless, it is even more challenging for general radiologists who are not specialized in neuroradiology, who may in fact be the ones reading these scans in the majority of cases. Therefore, there is an unmet clinical need to improve the review of these patient images, since most of them are reported by general practitioners. This project seeks to leverage Toronto Western and the University

Health Network's vast and existing database of CTA and MRA scans, and the expertise in vascular conditions found at these sites, in order to implement machine-learning approaches for the detection and localization of brain aneurysms.

The goal of this project is to build an automated detection and localization methodology for brain aneurysms. First, to create a brain aneurysm dataset, a scalable platform will be built for a more streamlined workflow of precise, per-slice multi-annotation. Once established, this dataset will be exploited for model training and evaluation. To begin, baseline models using pre-existing techniques of 2D convolutional neural networks and maximal intensity projections will be created. After this, models will be trained using 3D convolutional and fully convolutional neural networks, and 3D U-Nets, in a two-stage architecture for the screening and subsequent discrimination of aneurysms.

AS: Marzyeh Ghassemi

IS: Vitor Mendes Pereira

Learning More with Less

Sneha Desai

At bridge7 Oncology, we are developing an AI solution that increases the quality, efficiency and consistency of cancer treatment plans in radiation therapy. Our pipeline trains one model per protocol (cancer type), using gold standard data, and uses them to score a new plan for a specific cancer type. But what happens when our gold standard database for a specific protocol doesn't have many plans? We could learn a model with a limited number

of training examples, however, it wouldn't perform too well. By using techniques in meta-learning, we can learn attributes from many different cancer types to initialize a model in such a way, that it will only require a limited number of training examples to learn a protocol that has a small amount of data. We are currently evaluating this approach using generated data, then we will transition to evaluating clinical data.

AS: Marzyeh Ghassemi
IS: Chris McIntosh

A Solution Focusing on Mapping and Analyzing Customer Journeys

Zhou Fang

A customer's journey involves all of their interactions with a company, product or service. Analyzing customer journeys allows organizations to understand their markets, and the ways customers interact with them.

insights. Data integration projects focus on unifying data without providing customer insights.

- Lack of real-time customer analytics. Traditional analytical approaches provide a historical statistical view.

The traditional methods of journey mapping has serious limitations. These are as follows:

- Can only deal with specific data sources.
- Data integration is difficult. Data exists across different systems and data integration typically take months.
- There is a gap between data and

In this work, we developed a generic solution for building the data warehouse for different data sources. By using the latest techniques in machine learning, we analyzed customer behavior, allowing us to help companies prevent churn, measure the similarities between customers and provide recommendations for the next best actions.

AS: Nathan Taback
IS: Periklis Andritsos

Clinical Decision Support: An Explainable Language Model

Jinyue Feng

Free-text medical notes contain predictive information not captured by structured data. Deep learning models without explainability are often met with resistance in medical contexts. Sitting at the intersection of explainable AI and clinical language

processing, this project aims at constructing a multi-task language model that also provides rationales for its decisions. In the future, this model will be applied to other language-related tasks and extended as part of an explainable decision support system.

AS: Marzyeh Ghassemi
IS: Frank Rudzicz

Learning from Mistakes: An Active Learning Framework

Jinyue Feng

Most of our data coming from the OR Black Box® takes the form of hours-long videos, which require a huge amount of time to annotate. Active learning aims to reduce such effort by selecting the most informative data points and improving models with fewer labels. Training objective functions are crucial for any neural networks, and the loss values contain both model-related and data-related information. The general goal

of this active learning method is to train an observer model to predict the performance of the model of interest on unlabeled data, given its behaviour on labeled training data. The samples with the largest estimated loss values are selected for labeling and expected to induce the greatest model changes. Our approach outcompetes other active learning methods on image classification tasks.

AS: Marzyeh Ghassemi
IS: Frank Rudzicz

Neural Networks Quantization and Pruning Techniques

Yingying Fu

Untether AI is building a power-efficient hardware accelerator for neural network inference applications. It has designed a custom Application-Specific Integrated Circuits (ASIC) that will run neural networks faster, while consuming significantly less power. The power efficiency is from the unique near-memory computing architecture, where large arrays of small processing elements are placed close to memory. This saves power by reducing data movement between memory and compute. This accelerator is also optimized for performing low precision integer arithmetic.

The goal of this project is to explore the quantization and pruning techniques of

AS: Gennady Pekhimenko
IS: Darrick Wiebe

neural networks. Since the accelerator is built for handling integers, model quantization is an important step. Errors introduced from quantization can impact the accuracy of the model's predictions. We explore different quantization techniques that work most efficiently on the hardware, while maintaining accuracy. In addition, we study different pruning techniques to change the network into a sparse network, which can result in additional power-saving on the hardware. Overall, the emphasis is on co-optimizing the unique hardware design and transforming state-of-the-art neural networks to run more efficiently on the hardware.

Large-scale Semantic Textual Similarity (re-)ranker with ElasticSearch and Universal Sentence Encoder

Zihang Fu

In an automated FAQ system, the system takes as input a query question asked in natural language, and retrieves the most similar question with the corresponding answer from the database. In this project, we improved a large-scale FAQ system, based on semantic textual similarity matching. We considered the problem of retrieving a semantically similar question from a very large database as a two-stage problem. The first stage was using a fast but coarse algorithm to filter out the top 1000 questions (candidates) from the

AS: Graeme Hirst
IS: Luchen Tan

database, that are possibly similar to the user's query. The second stage is to use a slower but more accurate NLP model to find the best candidate. Compared with only using the NLP model in the second stage, the formulation of the problem greatly reduces the time cost with very little drop in accuracy. We run experiments with multiple NLP models and find the Universal Sentence Encoder works best.

On the Road to Joint Perception and Prediction

Stephen Gou

A key capability of self-driving vehicles is perceiving the world, including detecting objects, tracking objects over time, and predicting the objects' motion. The classical approach involves independent models for detection, tracking, and prediction. However, there are several drawbacks, including the fact that it's time consuming to train and deploy models. Uber ATG published the first paper that presented the idea of performing detection, tracking and prediction jointly in a single model in the ground-breaking paper, Fast and Furious [1].

However, through empirical experiments and testing, it has been observed that predicted vehicle trajectories cannot follow the road very well, especially on roads that are curved. The goal of this project is to improve the model's ability to use the rich lane graph data in and HD map to improve predicted trajectories. This model extracts map feature

AS: Murat A. Erdogdu
IS: Inmar Givoni

per detection from the detector and then applies a refinement network to generate new trajectory predictions.

Per detection map feature is computed through the use of lane graph data. The resulting feature is a rasterized image of the current driving lane ahead of the actor. Each rasterized local map feature will first be transformed into an actor's local space, before going through a U-Net block to extract embedding. The following fully-connected layers will output the new trajectory predictions, which will be transformed back to the ego vehicle's space.

Through experiments using Uber ATG's proprietary TOR4D dataset, this showed that this method reduced the L2 errors of predicted trajectories, especially on curved roads.

Bandwidth-Efficient Transaction Relay in High-Throughput Blockchain Systems

Yilin Han

In actually deployed blockchain systems, transaction dissemination is a critical process to ensure transactions are being packed into blocks in time, and this process consumes a significant portion of the available network bandwidth. Prior work investigated the network consumption issue of transaction dissemination in classical blockchain systems like Bitcoin. While this issue is much more important and challenging to handle in the situation with

more advanced high-throughput blockchain systems, such as Conflux and OHIE where network bandwidth is already the bottleneck of the system performance.

In this project, we propose ShRecon, a novel transaction relay protocol designed for optimizing bandwidth efficiency of high-throughput blockchain systems, where network bandwidth is the bottleneck and scarce resource.

AS: Sam Toueg
IS: Fan Long

Machine Learning Forecasting Techniques For Company Financial Fundamentals In Long-Term Value Investing

Thomas Hollis

As high-frequency trading becomes an increasingly popular and saturated problem domain, Valsys seeks instead to explore a somewhat neglected frontier of alpha-generation within the field of quantitative research: long-term value investing. One of the most commonly used financial models for forecasting a company's intrinsic value is the Discounted Cash Flow (DCF) model. This model relies on the projection in perpetuity of a company's cash flows, to derive a total sum of the company's forecasted profits. This sum is then discounted back to today to reflect both the time value of money and the uncertainty of the cash flow forecasts. This output valuation of the firm is then divided by the total number of shares outstanding, to generate an implied share price based on the company's fundamental value.

AS: Roger Grosse
IS: James Worthington

This project aims to improve the DCF modelling process by leveraging various machine learning models to improve forecasts, interpretability and analyst feedback. The DCF approach to valuation can be framed as a semi-supervised learning problem, since we have labels for each forecast of company fundamentals but no true label for the final DCF output valuation of the firm. Recurrent Neural Networks (RNNs) have increasingly shown promise as deep learning frameworks for financial time series (FTS) forecasting. To help alleviate the vanishing gradient problem of vanilla RNNs in multivariate multi-step ahead forecasts, LSTMs can be leveraged, along with a variety of other specialised modifications, to improve on current FTS benchmarks. In addition, feature importance, which can be extracted from Random Forest ensembles, can be used to help analysts determine line items likely to contain the most signal.

Investigation on Hand Pose Estimation and Tracking in Interactive Space

Jing Huang

As one of the most important kinds of information facilitating human-computer interaction, human action, gestures, and others have received much attention, within both academic research and industrial development. The development of low-cost consumer depth cameras, and the success of deep convolutional neural networks in all the computer vision tasks on pose estimation, have made human pose estimation and tracking and gesture recognition trend in the computer vision field. Many valuable applications and ease-to-use services can be enabled, based on this technology. Based on different quality requirements of the services and the technology readiness level, many

AS: Arvind Gupta & Sanja Fidler
IS: Jeff Gorda

approaches have been developed.

The objectives of this research are to develop deep learning and computer vision-based technologies to solve the challenging problems of human pose and hand pose estimation and tracking, to provide the user with an interactive visual experience. We are aiming to develop the problem-driven solution, and investigate the feasibility and solution of integrating visual features and depth information, within the component of a deep neural network to boost the recognition performance. The performance will be evaluated both quantitatively and qualitatively using real problem domain data to assess the impact of this research.

Enabling Purchase of Residential Homes at Scale

Bharadwaj Janarthanan

Institutional Buyers (iBuyers) ease the process of selling by enabling the quick purchase of homes at market prices, and thereafter owning the selling process of the home to a prospective buyer. There are significant capital risks involved in buying and selling homes directly related to the time they spend on the market. Current methods for liquidity assessment are predominately informed based on market trend aggregates and operational expertise,

AS: Nathan Taback
IS: Craig Dunk

making it less desirable as a scalable process. The key goal of this project is to use state-of-the-art survival analysis techniques at discrete time intervals to predict, with good precision, time on the market for listed homes to sell. This enables advancements in the art of marketplace modeling beyond predictions of clearing price, into more complicated time domains that are relevant for asset classes, that have a carrying cost.

The Right Initialization for Transformers

Xiaoshi Huang

We studied the component of transformer models, analyzing the gradient updates through attention mechanism. From this we identify an initialization scheme that eliminates the need for both warmup and layer normalization in transformer models, as well as propose an improvement to the attention module that boosts model performance.

AS: Jimmy Ba
IS: Felipe Perez

Explaining Cancer Treatment Quality

Angad Kalra

Radiation therapy is an expensive, timely, and inconsistent practice. The quality assurance of a typical treatment plan takes an average of six hours and costs upwards of \$6000. bridge7 Oncology has developed an AI platform that improves the efficiency, consistency, and quality of cancer treatment plans. Institutions and medical clinics around the world are now able to improve their treatment protocols and deliver high-quality care, even if medical expertise is not readily available. However,

black-box algorithms are not trusted nor safe in healthcare, which is why it is crucial that our software provides evidence for the decisions it makes. The purpose of this project is to improve the explainability of our AI platform in ways that are easy to comprehend by our clients. This includes experimenting with the latest research, developing novel methods, and investigating different User Interface (UI) approaches that can help our users quickly determine how they should improve their treatment.

AS: David Fleet

IS: Chris McIntosh

Dynamic Pricing for Optimizing Demand and Profitability

Jing Li

Price optimization is an active area of research and development in the retail industry. It is a critical lever for optimizing both demand and profitability. In particular, the partner organization needs to match the premium charged for “PC Express” orders to demand. The most immediate motivating factor has to do with order pickup timeslots: popular timeslots are congested, making it harder to deliver a positive customer experience; while the other windows are idle and hence overstaffed. To this end, the partner organization hopes to smooth the demand curve, which should improve both labor efficiency and customer experience. In the project, large volumes of purchase

history data including pickup and delivery “hot-spots”, historical pricing information, historical inventory, and unfulfilled orders were to be utilized to develop the dynamic pricing model. During the initial stage of the project, an offline model that maps time-of-day and day-of-week to ideal price points was developed. Once the initial objective was achieved, the project was carried out to develop an online system that could react to demand spikes, with consideration of external data sources such as weather and news. An article level dynamic pricing model was explored, along with the main project, to drive both demand and net-margin.

AS: Roger Grosse

IS: Richard Downe

Active Vision in Robotics for Optimal Mobile Camera Control

Angran Li

In the bin-picking problem, poor Object Detection and Pose Estimation ODPE results can be caused by reflection spots on objects, due to lighting conditions and their surface material. This project uses a mobile camera to solve this problem: the camera moves to different viewpoints, fuses captured data and complements the missing parts. Since the task is infeasible to be solved by hard-coded rules and manually labeled data, the reinforcement training method is used for learning a sensible policy.

To solve the real challenge of fitting a reinforcement learning algorithm to real-world robotics problems, first the training environment is carefully defined to incorporate trajectory memories in the Partially Observable Markov Decision Process (POMDP) environment. Second, the training data is sampled offline and the

policy is updated online, for fast iteration and to avoid the disadvantage of slow convergence.

The biggest problem comes from depth maps taken by the mobile camera, which serves as the observation of the RL agent. The depth information is noisy. Hence the policy converges to a situation that for all the different observations the agent tends to, executes the same action. Several methods are tested for potential solutions: autoencoder, which introduces a huge training cost with scattering reflection spots; point cloud features, which are not representative enough for the global features of different camera poses; and a heatmap of point pair feature (chosen as the observation representation for the current algorithm), which takes the confidence of pose estimations and infers the area that needs more attention from the mobile camera.

AS: Florian Shkurti

IS: Dan Rico

Smoothed LIME for Time Series Prediction

Casey Li

This project investigates a potential solution for the commonly observed numerical instability of Local Interpretable Model-Agnostic Explanations (LIME) when used as an explainability technique on models that output predictions for the same entity over nearby time-points. Given an unexplainable model F , LIME proposes to fit an explainable model G (commonly, a ridge regression) which has two properties, represented in the loss function used to train F : 1) local

fidelity - as in, behaves similarly to F around a given point; and 2) is relatively less complex. This project proposes the addition of a third term to the loss function, which explicitly rewards solutions that are stable for similar points over time, and provides insight on the practical and computational challenges of implementing such a solution in a production environment to produce client-facing output.

AS: Nathan Taback

IS: Erik Pazos

A Facial Sign Scoring Model and DNN Compression Application

Zeqi Li

Convolutional neural networks (CNN) such as ResNet, UNet and MobileNet excel in various general computer vision tasks, and the area of medical diagnosis benefits significantly from applying those models. In this work, we developed a CNN model to automatically grade the severity of multiple facial skin signs through 'selfies' pictures, taken by women of different ages and ethnics 'in the wild'. The model leads to global agreements with the expert's assessment and offers a new quantitative

AS: Sanja Fidler & Arvind Gupta
IS: Irina Kezele

approach in the quantified description of facial skin signs.

To further reduce the size and lower the computation requirement of the model, we explored several parameter pruning algorithms for doing model compression. For models like ResNet50, we are able to reduce the model size by 30-40% and FLOPs by 40-50% without degrading the original model performance.

Multi-person 3D Poses Estimation with RGB-D Cameras

Mengxuan Lyu

Human poses convey rich information, which may assist in multiple computer vision tasks including action recognition, people re-identification, and industrial robots. Therefore, understanding human poses in the Operating Room (OR) would be useful for surgical activity recognition and surgical skill analysis.

AS: Arvind Gupta & Sanja Fidler
IS: Frank Rudzicz

Moreover, due to the OR's strict requirements for sterilization, Lynch et al. (2009) found that increased traffic in the OR would increase the contamination around the wound, leading to a high risk of potential infections. Consequently, estimating human poses and tracking people would help to understand certain aspects of surgical procedures and improve surgical safeties.

Cervical Spine Fracture Detection

Pulkit Mathur

A significant number of trauma patient present to the emergency room with cervical spine fractures each year, with delays in the operative measures often resulting in irreversible damage (i.e. paralysis) which may increase the risk of death. Immediate fracture detection in the ER room could help physicians to make a swift decision as to whether the patient requires surgical intervention, without waiting for the report from radiology. We obtained a total of 729 positive and 3029 negative cases that included cervical spine computed tomography (CT) scans of patients from the St. Michael's database. Each case comprised of a sequence of images of a given patient's cervical spine, starting from the top vertebra C1 to the lowest vertebra

AS: Shahrokh Valaee
IS: Hojjat Salehinejad

C7 and are annotated at the pixel level with the help of three radiologists and verified by a board-certified neuroradiologist.

Using this data we trained an object detection model and integrated the model results with an algorithm that uses the sequential information of images to reduce the false positives. We later compared the model results with radiologists and neurosurgeons. The results showed that the model was on par with humans, and better in some situations. We aim to implement this approach in hospitals as, not only will it help in reducing human errors, but it will also reduce the time between registration of the CT scans and its diagnosis by a doctor from an hour (on average) to a few seconds.

Detecting Brand-Specific Purchase Evidence on Twitter

So Hyun Park

Social forums constitute powerful media through which people can express their needs and opinions on a daily basis. Therefore, applying Natural Language Processing (NLP) technology to the sheer volume of social media data offers tremendous opportunities for businesses. For example, Purchase Evidence (PE) has been shown to predict sales growth in publicly traded companies [1]. Therefore, detecting PE on Twitter can help generate useful trading strategies in the stock market.

AS: Yang Xu
IS: Hugo Mailhot & Mickael Temporão

Here, PE is an implicit or explicit sign of prior purchase or imminent indication of purchasing a product or service. "That new Georgia Gold chicken at KFC is actually really good" exhibits PE explicitly, whereas "A girl in the @netflix show #Glitch has my birthday" does not exhibit PE. The purpose of this study is to build an effective NLP model that automatically detects PE, given a tweet that is relevant to a particular brand.

Table Extraction from PDF documents

Chenzi Qie

The Management Information Circular (MIC) is a document for a company's shareholders, outlining important matters on the agenda at the annual shareholders' meeting, or at a special shareholders' meeting. The MIC documents are typically issued on an annual basis on behalf of the company. Financial institutes, like Scotiabank, are interested in staying up-to-date with shareholder information, including their shareholder's names, number of voting shares they hold, company stock symbol, etc., which is included in the MIC documents. Previously, this information was manually collected by reading the PDFs or purchased from third parties. There are thousands of documents every year and some have hundreds of pages. It was an extremely time-consuming task before the automatic table extraction

algorithm was developed.

The ultimate goal of this project is to build a tool to extract the shareholders' information from tables embedded in PDF documents. The difficult part is that MICs issued by different companies tend to have different table formats to store the shareholders' information. Despite this, the current algorithm has successfully extracted the shareholders' information from target documents. The results from past documents are saved in a structured table and can be easily queried. The algorithm is processing coming MIC documents on a daily basis, bringing great value and convenience.

AS: Gerald Penn
IS: Simrandeep Singh

Quantitative Analysis of Contact Dependent Interactions Between Cancer Cells and Fibroblasts

Sumeet Ranka

Cell-cell interactions have been found to have a significant impact on drug efficacy. Recent literature has shown that the tumour stroma, the connective tissue surrounding cancer cells, can promote the resistance of cancer cells to many treatments. In this project, the following goals are achieved:

1. Multi-cell cultures are segmented and cellular features are extracted using CellProfiler. Different cell-types (cancer vs.

fibroblasts) and cell-type boundaries are classified based on these features.

2. Interactions between cancer cells and fibroblasts are characterized using classical features like intensity, shape, and texture after identifying cell-types that are in contact with each other.

3. The pipeline is automated to be deployed to screens performed week after week.

AS: Jimmy Ba
IS: Oren Kraus

DNN for HS Code Classification

Andrejs Ruzencevs

The Harmonized Commodity Description and Coding System, which is also known as a Harmonized System (HS), is an internationally standardized system of names and numbers to classify traded products. The Canadian Harmonized System contains more than 10 thousand trade codes, where each code is represented by 10 digits.

Fortunately, it was discovered that complex deep neural networks mixed with natural language processing can give quite phenomenal results. Given a relatively short description, accompanied by additional features such as customer and vendor codes, the deep neural networks can provide a meaningful and accurate output.

AS: Arvind Gupta
IS: Reddy Nellipudi

Multi-Modal Representation Learning for Video Captioning using Variational RNNs

Vaibhav Saxena

Generating linguistic descriptions for videos has a multitude of potential applications in a variety of fields. For Surgical Safety Technologies, automatic caption generation can automate a number of tasks for the analysts who annotate videos of surgeries and operating rooms. Through automatically generated captions, we can build a search system that can be used to find videos of interest in the surgical video database. In this project, we focus on the task of video captioning using dense embeddings

as obtained from a Variational Recurrent Neural Network (VRNN). VRNNs provide a probabilistic model for inferring latent variables, which have proven useful for modeling handwritten sequences and speech. This project explores the usability of VRNNs for modeling videos, and further use of the inferred latent space for caption generation.

AS: Jimmy Ba
IS: Frank Rudzicz

Communication in the Operating Room

Chantal Shaib

In high-pressure environments such as operating rooms, unimpeded communication is imperative to the procedure at hand. Analyzing annotated dialogue during a procedure can provide important information that can be used in downstream tasks, such as sentiment analysis and uncertainty detection. This

project focuses on defining uncertainty and creating a library that parses both audio and text, according to both learned and predetermined features. These features are used to detect and score uncertainty in the OR.

AS: Graeme Hirst
IS: Frank Rudzicz

Identifying Subtle Phenotypes from High-Content Imaging

Srinivasan Sivanandan & Jeevaa Velayutham

High-throughput screening assays, along with high-content imaging technologies, enable rapid and large scale measurement of single-cell phenotypes for screening small molecules for drug discovery. However, the expressed phenotypes vary with changes

in the small molecule concentration, with the phenotypes being subtle at lower concentrations. In this work, we have developed a new feature extractor using loss functions that are robust in classifying cells at lower small molecule concentrations.

AS: Arvind Gupta & Sanja Fidler

IS: Muneeb Sultan & Hervé Marie-Nelly

Identity Resolution in Digital Advertising

Kavita Srinivasan

Advertising is as much about knowing when not to promote a product as it is about when to. For example, when running alcohol and cannabis ads, the advertiser must ascertain that the user is aged 19 or above. Nowadays, people utilize multiple devices and channels to complete various tasks - they may check the weather app on their phone, do some shopping on their tablet, and research work material on their laptop browser. Age information may only be available on a channel different than the one through which the user is connecting. This fragmentation of online identities makes it difficult to gain a holistic understanding of users across channels.

Pelmorex Audience, the mobile advertising division of Pelmorex Corp., would like to

AS: Marsha Chechik

IS: Lubna Khader

create a master set of user profiles that identifies different representations of the same user across channels. This will enable the company to enhance the user experience by limiting ad repetition on different environments, whilst also customizing ads to the interests of users as observed across devices. Pelmorex Audience does not request personally identifiable information (PII) to provide access to its services and, therefore, does not receive or store PII data - users' representations are anonymized. The outcomes of this research will contribute to the active research area of large-scale identity resolution. Advanced probabilistic matching techniques, such as indexing, attribute-level comparison, and match classification will be explored.

Distributed Schema Registry with Confluent Kafka

Duc Truong

Transactional data at Ethoca requires a global scaled management system, which has high availability, durability, read and write performance. Besides that, the system must ensure, not only the confidentiality of the sensitive data stored on the system, but also the seamlessness of data transfer across different regions. To satisfy the strict requirements, global service providers usually set up their applications and backend databases on multiple data centers in different geographic regions. Schema registry is brought to the picture as part

of the solution, to increase the efficiency and consistency of applications and stream processing pipelines. This project seeks to determine a reasonable path forward in technology design and adoption, to accommodate current and anticipated data privacy and data residency concerns, making it possible to keep data within borders while still facilitating collaboration across those borders.

AS: Matt Medland

IS: Dale Woolridge & Hrant Arakelyan

Improving a Smart Agriculture Solution

Antoine Viscardi

With the goal of making agriculture more sustainable, and fuelled by breakthroughs in IoT technologies, the development of smart agriculture solutions has been booming in recent years. They give farmers access to a quality and a quantity of data never seen before. However, it often leads to confusion and a feeling of overwhelmedness. To understand, imagine that you were not being told that there was a 83% chance of rain today, but instead were given a collection of meteorological maps and atmospheric measurements, and you now had to decide if you should bring an umbrella to work. Ukko Agro aims at solving

this problem by providing farmers with actionable recommendations, as opposed to abstract numbers and reports. This requires recommendations to be delivered in real-time, which involves streaming field-level readings to cloud-based disease forecasting models. It is known that a key predictor of the spread of crop diseases is the relative humidity level. Forecasting this predictor at a hyper-local level could enhance the current disease forecasting model's accuracy. Hence, this is what this project aimed to do.

AS: Scott Sanner

IS: Avi Bhargava

Personalizing Energy Savings for Smart Thermostat Users

Laura Walsh

ecobee is a Canadian hardware and software company that brought the first smart thermostat to market in 2009. Since then, it has released increasingly capable thermostats and has expanded its hardware offerings to include remote temperature/occupancy sensors and light switches. In addition to allowing a user to monitor and adjust their home temperature remotely, ecobee smart thermostats offer a suite of automated energy saving features that intelligently reduce energy usage without compromising user comfort.

This project is concerned with using machine learning to personalize and optimize one such energy saving feature: the Demand Response Event. A demand response event is a period of time in which a

AS: Maryam Mehri Dehnavi
IS: Inayat Khosla

utility provider asks participants to use less energy than they normally would, to offset a spike in energy demand. Without this program, utility providers have traditionally had to either buy energy at a high premium, or utilize peaker plants (costly, non-eco-friendly power generators) to meet the demand and prevent blackouts. For this project, historical data on user behavior and event characteristics were used to train a classifier to predict which users would stay in these events. The best performance was achieved using the extreme gradient boosting algorithm. This information can be used to tailor the length and depth of the events to each user's preferences. Maximizing participation in these events whilst minimizing user discomfort are

Fair and Robust Treatment Effect Estimates: Estimation Under Treatment and Outcome Disparity with Deep Neural Models

Shirly Wang

Treatment effect estimation is a statistical tool often used to compute the efficacy of a treatment in the presence of other factors, or confounders. However, confounders are often manually chosen by a clinician or domain expert without rigorous validation, and may not take protected identities (denote as A) - or known causes of treatment/outcome inequalities into account. We explore the formulation of fair and robust treatment effect estimates under various potential disparities using deep neural models. We use a semi-synthetic dataset under four generative settings, 1) A is unrelated to treatment and outcome, 2) A only causes treatment disparity, 3) A

only causes outcome disparity, 4) A is an observed confounder.

We find that deep neural models provide more accurate treatment effect estimates, and that including A as a covariate helps in learning more reliable and fair causal effect estimates in all cases except 2) exclusive treatment disparity. Our findings suggest that manually selecting covariates, as is commonly done for propensity scoring based treatment effect estimation can lead to inadvertent biases (by virtue of excluding the source of disparity A) on effect estimates.

AS: Marzyeh Ghassemi
IS: Maks Volkovs

Assessing and Addressing Health Disparities Related to Utilization of Preventive Care Services in Ontario

Shirly Wang

Health disparities are the avoidable differences in health outcomes at a population level, as a result of certain sub-populations facing social, economic, or environmental disadvantages. In an age where machine learning models become more popular in the healthcare sector, health disparities present new challenges. Preventive care services, such as screening and preventive medicine, are an essential component of a well-functioning healthcare system. As they allow for early diagnosis

and timely interventions, disparities in their utilization present challenges to medical fairness. Building on a thorough analysis on preventive care utilization disparities and their impact on medical fairness, this project aims to utilize the potential of deep learning models in predicting the needs for preventive care services. In addition to technical capability, the model will be trained with awareness of current health disparities and evaluated against fairness measures.

AS: Marzyeh Ghassemi
IS: Maks Volkovs

GANs for the Beauty Industry

Haicheng Wang

As one of the leading companies in beauty AI, ModiFace commits itself to combining cutting-edge technologies with real-world Augmented Reality (AR) applications. Our products range from the virtual try-on for brands from L'Oreal, to highly customized AR and diagnosis applications driven by deep learning. Generative Adversarial Networks (GANs) has been a heated

research field for half a decade, leading the performance for image generation, style transfer and various other applications. We are particularly interested in GANs because of its high capacity for realistic, high-resolution image generation. Through this project, we will share our exploration and experience of applying GANs to the beauty industry.

AS: Arvind Gupta & Sanja Fidler
IS: Irina Kezele

Multi-Modal Representation Learning

Lan Xiao

The world around us is complex, and as such we experience and interact with the world using multiple modalities. We see objects, track their movement, and recognize actions involving those objects; we hear and interpret sound and speech; and we use spoken language to enquire about the environment. In order for an AI system to interact with the world in the way we do, it must be capable of seamlessly joining the various sources of information (a.k.a.,

modalities). The AI research community has recently paid much attention to building multimodal representations and systems. Examples include: cross-modal (document, image, video) retrieval, dialogue-driven video search and summarization, image/video captioning, visual question answering and visual dialogue. This project focuses on building models that enable natural user—device interactions, by building multimodal representations and models.

AS: Arvind Gupta & Sanja Fidler
IS: Afsaneh Fazly

Real-time Axis-based SLAM in Geotechnical Engineering

John Xu

RockMass is a mining technology startup improving operational efficiency and safety. Every day, geotechnical engineers estimate the orientations of planar surfaces in rock masses to analyze a mine's stability. The most widely used approach to collect these estimates uses hand measurements with a compass and inclinometer. By leveraging the technique of 3D axis mapping, it is possible to automatically transform data collected from a moving platform into these orientation estimates. The speed and efficiency of collecting data are drastically improved compared to manual techniques and stationary LiDAR.

The proposed research seeks to develop a lightweight axis-mapping SLAM algorithm into a real-time system, and

AS: Alec Jacobson
IS: Stuart Bourne

extend its capabilities to produce a novel data collection tool for geotechnical engineers. The first step in the research is to understand the axis mapping research landscape and analyze the computational bottlenecks of the process. A proof-of-concept posteriori-based SLAM design, written in a 2016 doctoral thesis, is materialized and then validated against manually verified observations in real mine sites. The prototype algorithm is then adapted to execute in real-time on an integrated embedded platform. The goal is to produce rock orientation estimates at industry-leading levels of speed, accuracy, and functionality. Key optimizations include modern GPU parallelization and point cloud methods such as the Maximum Consistency Minimum Distance algorithm.

Advanced Analytics in Multiple Sclerosis Research

Zhen Yang

Multiple Sclerosis (MS) is a chronic neurological disease affecting the central nervous system, and is the leading cause of neurological disability in young adults. The MS clinic at St. Michael's Hospital (SMH) is among the largest in the world – it is staffed by eight MS neurology specialists and accommodates over 7,000 MS patients. While considerable data are collected in the form of structured (e.g., lab test results) and unstructured data (e.g. digital text notes), the ability to glean this information to assess the quality of care and conduct advanced analytics are limited. Leadership at the MS clinic at SMH and the LKS-CHART propose to work together to advance two critical areas of clinical care and research: 1) development of a research database based on automation of clinical information extraction and analysis and 2)

AS: Marzyeh Ghassemi
IS: Muhammad Mamdani

predictive modeling using existing clinical data to predict disability progression and patient outcomes.

The aim of this project is to create a pipeline to extract various clinical variables using Convolutional Neural Networks (CNNs) trained on top of our own word embeddings (Word2Vec), based on MS consult notes with the objective of automating clinical information extraction and reporting for the MS clinic at St. Michael's Hospital. Current progress shows that we can successfully extract various clinical variables (e.g. Expanded Disability Status Scale scores) using a multi-model approach. Future work will be focused on using preliminary results to develop advanced models to predict important patient outcomes, such as the risk of patients' disability progression.

Audience Allocation to Automatically Detect Retail Geo-clusters using K-MEANS, DBSCAN

Emily Yang & Anxin Zhao

Pelmorex delivers information and data solutions to clients throughout Canada by operating various media channels, such as The Weather Network and MeteoMedia, and their associated digital properties. One of these digital properties is The Weather Network application, one of the top-rated and most downloaded software, developed for both iOS and Android smartphones/tablets. One of the key features of the app is the ability to provide local-weather information, as well as the 'severe weather' alerts once a user of the app has given consent for the app to access the device's geo-location data. All data is collected in compliance with existing Canadian privacy laws.

The goal of the research project is to allocate geo-location, dwell time, the number of visits and other attributes of each audience

AS: Nick Koudas
IS: Marek Kos

member to a retail cluster. An initial clustering is accomplished by running a Google Dataflow. The initial clustering uses windowing, based on the distance between sorted location breadcrumbs and is able to ingest around 20 billion device location breadcrumbs. Then an unsupervised machine learning clustering algorithm is developed, by utilizing K-means and DBSCAN. This algorithm is also designed for parallel running at locations across Canada and, due to the large data size, PySpark and BigQuery are applied in this phase. Finally, we generalize the probability distribution for each cluster and assign a z-score to each visit. As a result, visits are continuously updated to each cluster daily. This project represents a breakthrough since it enables recognition of location patterns, which in turn results in an improved user and client experience.

Designing a Machine Learning Framework to Study the effect of Polypharmacy on Seniors with Multimorbidity

Seung Eun Yi

Despite a significant improvement in healthcare systems over the past decades, the rapid growth in the number of patients with multiple chronic diseases - called multimorbidity - stands as a complex challenge to healthcare services primarily designed to treat individuals with single conditions. In 2013, 65% of Ontario residents died with more than five chronic illnesses, representing an increase of 39% since 1994. Alongside multimorbidity, the number of people exposed to polypharmacy, a term referring to concurrent use of five or more drugs, has been increasing as well. Sixty-six percent of seniors in Canada are prescribed with five or more different drugs, and 26.5% with ten or more. However, half of the medications given to seniors are either prescribed in a redundant quantity, or not properly taken. This suggests that treatment

AS: Marzyeh Ghassemi
IS: Maks Volkov

plans are suboptimal and lead to an over-medication of the population, exposing patients to dangerous polypharmacy side effects.

The goal of this project is to study polypharmacy effects on seniors (aged 65+) with multimorbidity. The first step is to build a variety of machine learning and deep learning models to predict conditions leading to avoidable hospitalization (conditions for which hospital admission could be prevented by interventions in primary care) of seniors, given their conditions and medications. This will help us understand which combinations of medications can lead to avoidable hospitalization and allow us to analyze how this hospitalization could be prevented by modifying sets of medications.

Page 29 - ARIA 2019

Speaker Diarization with Time Delay Neural Network

Ziming Yin

Speech recognition has been around for decades, but with the emergence of deep learning speech recognition it is now much more accurate and reliable. Two years ago, Scotiabank developed and built an end-to-end speech-to-text transcription, topic modeling and sentiment analysis system to transcribe audio calls, in order to identify key topics and associated segment level sentiments of the calls. However, the transcription quality could be improved if we upgraded the models by leveraging the knowledge of deep learning.

For this project, we started with the first step of our end-to-end pipeline - diarization - which essentially segmented the call into

AS: Gerald Penn
IS: Saif Charaniya

A Study on Model Deployment: From Research to Production

Zhening Zhang

In recent years, the extensive development of neural networks, such as ResNet[1], MobileNetV3[2], have achieved an impressive performance on computer vision tasks. At the same time, a number of frameworks are readily available for developers and researchers, allowing them to model and train their own Deep Learning (DL) systems within their applications. However, two problems occur when deploying models to production:

- 1) Commonly used architectures are

AS: Sven Dickinson
IS: Irina Kezele

sentences, and then clustered sentences into different speakers. Previously, the output segments of the diarization model were mostly short pieces, which resulted in low performance in the following language model. We have now replaced the model with a time delay neural network-based framework to Scotiabank datasets, by domain adaption and model optimization. In the first stage, we conducted several experiments on different models to compare performance. We then fine-tuned the model with the public speech dataset, as well as the Scotiabank dataset. The final diarization error rate decreased from the previous error rate by 42%, and the word error rate decreased from the previous one by 15%.

often large in size and have hardware requirements, which make it difficult for them to be deployed on real-time mobile/web applications.

- 2) Each framework has its own particular format, which makes interoperability and usage of code and models trained inside different frameworks difficult. In this work, we develop a conversion tool between different machine learning frameworks, and explore several approaches in model compression.

Forecasting Battery Health for Mobile Devices Using a Generalized Additive Model

Feixiong Zhang

SOTI Inc. is a global company that provides enterprise mobility management. Insight Agent is a product developed by SOTI, that helps SOTI's customers collect various specifications from mobile devices, such as battery level, voltage, current, and other metrics. This research project is conducted to accurately forecast the health of batteries in mobile devices using the metrics collected by Insight Agent. After exploring the data collected from hundreds of thousands of mobile devices, we extracted

AS: Murat A. Erdogdu
IS: Hossein Taghinejad

the Key Performance Indicator (KPI) that accurately describes battery health. Using the estimated KPI, we explored many machine learning models for time series forecasting and proposed a generalized additive model that accurately predicted battery future health for mobile devices. This will help the customers minimize cost of device maintenance, as well as total cost of ownership, whilst ensuring a minimal impact and loss to business productivity.

Real VR Hands and Interaction with Virtual Objects

Shihang Zhu

Tactual Labs Co. is an innovator in human-to-computer sensing and processing technology, providing human-computer interaction in the domain of mixed reality, automotive and display. While Virtual Reality (VR) has adapted existing game console controllers, more expressive input devices are required to enable a more realistic and immersive user experience.

Tactual Labs's Game Controller enables real-time hand skeleton reconstruction and hand pose visualization, based on the capacitive sensor surrounding the free-formed surface of the controller. The usage of bare hands can improve the user experience in VR applications by providing natural, seamless, and immerse interactions. By enabling interaction between reconstructed hands and virtual objects, users can gain an unprecedented way to conduct social communication, productive tasks, and gaming.

AS: David Levin
IS: Ricardo Jota

The objective of this project is to transform the skeleton and pose data collected by the Tactual Controller into a standard input format used by existing VR gaming platforms, so that the controller can be used as a new way to interact with existing content. Currently, Steam is one of the most popular PC gaming platforms. It supports several consumer VR headsets, such as the HTC Vive and Oculus using its StreamVR tool. In this project, we transfer real-time pose data from the controller into SteamVR Actions, enabling full hand tracking in the VR environment. By achieving these features, users can customize gestures they prefer, to interact with virtual objects and play any VR games directly with Tactual Controllers in a more natural and lucid way. At the same time, these features help us to understand that the best way to apply the Tactual Controller on native applications is to control the process of controllers sending events.

MScAC Master of Science in Applied Computing



Computer Science
UNIVERSITY OF TORONTO



Mathematics
UNIVERSITY OF TORONTO



Statistical Sciences
UNIVERSITY OF TORONTO

Join the conversation online

#ARIA2019



@UofTCompSci

@UofTMathDept

@UofTStatSci

Master of Science in Applied Computing (MScAC)
Department of Computer Science
University of Toronto
40 St. George Street, Toronto ON M5S 2E4

mscac@cs.toronto.edu
mscac.utoronto.ca
+1 (416) 978-5180