



APPLIED RESEARCH IN ACTION

Master of Science in Applied Computing (MScAC)
2016–2017 Projects



Computer Science
UNIVERSITY OF TORONTO

MESSAGE FROM THE CHAIR ▶▶

Earlier this year, Canada's Fundamental Science Review, led by University of Toronto president emeritus, David Naylor, called on the Government of Canada to increase funding for fundamental research. Over time, new research discoveries can improve our health, make our cities smarter, and change the way we think about our world.

Applied research focuses on a shorter timeframe with more targeted goals. It examines a problem, uses evidence-based knowledge, and designs a solution. This is what our master's in applied computing students do best, as the projects here describe, and meshes perfectly with the always-evolving tech sector.

In computer science, we work on both applied and fundamental research, and welcome opportunities to partner with industry across our various master's and PhD programs. For further information on research partnerships, please contact our Associate Chair for Research, Professor Yashar Ganjali (acrir@cs.toronto.edu).

I hope you enjoy learning about this year's applied research projects.

Ravin Balakrishnan

Professor & Chair
Department of Computer Science



◀ MESSAGE FROM THE PROGRAM DIRECTOR

The Applied Research in Action (ARIA) showcase signals a major milestone for our 2016/17 cohort. It's not only the near end of the program's internship, but also marks an exciting new phase in their careers in technology.

We're proud to call the companies represented at ARIA our partners. You've seen the value and need for applied research in your organizations, and supported our students through hosting these exciting internships. Thank you.

Our next event will be the Applied Research Internship Expo (ARIE) where over three days our industry partners will have a chance to meet our 2017/18 cohort, who will start their internships this coming spring. If you are interested in hosting an applied research internship project, please contact us at msscac@cs.toronto.edu for further information.

Congratulations to our 2016/17 applied computing students!

Matt Medland (MScAC 2012)

Director of Professional Programs & External Relations
Assistant Professor, Teaching Stream

A Privacy Enhanced Mobile App for Decentralized Identity System

A federated identity ecosystem, in which relying parties (RPs) and identity providers (IdPs) have to connect only once, has many benefits. However, the principles upon which they are designed have several security and privacy gaps. The main problem arises from the fact that the IdP can be in a significant position of power, creating the potential to track an individual's transactions across multiple RPs and build up a large amount of identifiable information about them. The IdP has knowledge of user data, which is not directly required by the calling RP – this is a security concern and breaks privacy-preserving guidelines set by various governments. This project has involved enhancement of the security and privacy issue, guaranteeing the user's right to share their data without their activity being tracked by the IdPs or between RPs, by utilizing blockchain technology and distributed storage.

SecureKey is building a user-friendly mobile application for the iOS and Android platforms, that allows a user to register and authenticate with well-known trusted identity providers, such as banks or governments, and share identity information when a third party requires authentication. In our research, we aim to enhance the identity sharing system, so that the user can control their data, and release it to an authorized relying party, in the most secure and reliable ways. Our app will put the consumer back in the middle, enabling them to take control over their digital assets, deciding whether to share them and with whom, always with informed consent.

Md Obyadur Rahman, MScAC / **Industry Supervisor (IS):** Patrick Engel / **Academic Supervisor (AS):** Professor Marsha Chechik



An Artificial Agent for Light Switches

Ecobee is a home automation company that makes thermostats for residential and commercial use. The thermostats connect to the internet and help automate homes, while saving energy and money. Smart lights are the second-most desired home automation devices after thermostats, but being able to remotely control the lights is only a small part of the vision. Ecobee believes, by discreetly integrating intelligence into already existing devices like light switches throughout a home, it has an opportunity to extend its promise of whole-home comfort, control and savings without cluttering the home with more devices.

Ecobee

Ecobee light switches have embedded microphones, a speaker, and far-field voice technology that allows the user to control the lighting, thermostat, and other smart home products. One major part of this goal is to create an artificial agent, which monitors and controls the devices. This context-aware agent would observe through the occupancy, sound, illumination, and temperature sensors, and act upon the data acquired. To make this concept a reality the agent has to gradually learn the residents' schedule and behavioral patterns, and make decisions accordingly. Light switch agents are the first phase of this project, which aims to estimate the probability for turn on and off actions by leveraging machine learning and AI, in a non-deterministic environment.

Ramtin Rassoli, MScAC / **IS:** Sina Shahandeh / **AS:** Assistant Professor Roger Grosse



Automatic GIFI Mapping using Machine Learning

PwC deals with thousands of companies that all use different account number formats for their trial balances. A key challenge in any engagement is to accurately map each of our client's personal account codes to the GIFI (General Index of Financial Information) master list. This list is an index of items generally found on balance sheets, income statements, and statements of retained earnings. The Canada Revenue Agency (CRA) uses GIFI Codes to support their electronic filing and T2 processing systems.

The goal of this project is to develop an automatic GIFI mapping tool. The tool would use machine learning algorithms to process the descriptions associated with the client's accounts, to find matches with the GIFI account descriptions.

In order to find the most suitable model for this classification problem, six different models were tested: AdaBoost, decision tree, Gaussian naive Bayes, k-nearest neighbors (k-NN), Multi-Layer Perceptron (MLP), and Random Forest. 10-fold Cross-Validation was used to evaluate the models and Random Forest gave the most accurate predictions.

In addition, confidence scores generated along with the predictions were validated to establish a cutoff, that would further reduce errors and unnecessary work in practice, and save significant amount of resources spent manually determining and entering GIFI codes.

Tianyang Liu, MScAC / **IS:** Mike Maikawa / **AS:** Assistant Professor, Teaching Stream, Nathan Taback

Automatic Text Summarization of Financial Documents

Scotiabank has developed and built a data ecosystem that contains their proprietary data and data from third parties. This ecosystem contains data that has been transformed into network data structures and entity reference data. Scotiabank has combined these independent network data structures together and built a global entity management process, which takes data from the ecosystem and generates a unique entity. The definition of an entity is broad and may include retail, commercial, machines, financial institutions etc.

All public financial institutions are required to make their annual financial reports, management discussions, company decisions etc. publically available. These reports are highly unstructured.

Scotiabank would like to create an interface and a knowledge graph from the public data available. The interface would be used to query the public dataset, search financial reports from the dataset and extract important information from the financial reports. The extracted information will be used for future in-depth analysis.

Simrandeep Singh, MScAC / **IS:** Junjie Zhu / **AS:** Assistant Professor Frank Rudzicz

Automatic Speech Recognition for Cognitive Impairment Assessment

Speech and language deficits can serve as a key marker of neurodegenerative diseases, like Alzheimer's disease and Parkinson's disease, and aid in early detection. As these diseases progress, language functions can become affected across a range of linguistic and acoustic domains. Winterlight Labs is a Toronto-based company that is developing an AI diagnostic platform for detecting and monitoring symptoms of language impairment associated with dementia, aphasia and other cognitive and mental health conditions. The company's core technology extracts hundreds of features from snippets of elicited speech that can be used to measure lexical diversity, syntactic complexity, semantic content and articulation. Using machine learning and natural language processing algorithms, Winterlight can detect Alzheimer's disease symptoms with accuracy between 82% - 90%.

The objective of this research project was to construct an automatic speech recognition (ASR) pipeline to automate transcription, speaker diarization and utterance segmentation. The project focused on combining state-of-the-art deep neural network acoustic models with language models tailored specifically for cognitive impairment assessment tasks. Speech recognition was designed to be robust to reverberant and noisy environments. Decoding graph parameters was optimized using a dataset of elderly speakers diagnosed with Alzheimer's disease and 'healthy' controls. The ASR pipeline developed will expedite scoring and enable Winterlight to conduct further research into calibrating and refining its impairment classifiers.

Josh Ames, MScAC / **IS:** Liam Kaufman / **AS:** Professors Alex Mihailidis and Richard Zemel



Automation of Vehicle Insurance Claim Process

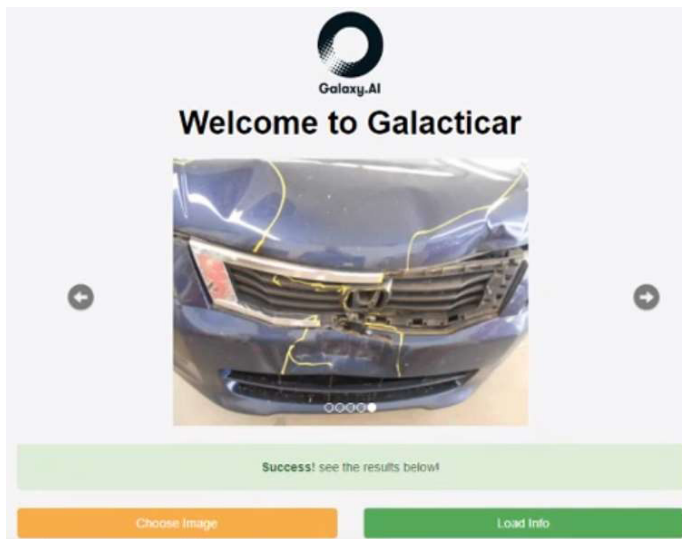
One of the fundamental tasks for a property and casualty insurance company is to assess the amount of damage that has occurred to a person's property, based on the images and information provided in a claim. This process is normally done by teams of human claim adjusters.

Galaxy AI has developed a new system, "Galacticar", which automatically processes images in a vehicle insurance claim and provides both an overview of the types of damage sustained to the car, as well as a price estimate for the amount of damage. The project was to build a system that would be able to extract various pieces of information from the images submitted in a claim, combine them with textual information to provide an accurate line by line breakdown of damage to the car, and convert that into a final price for repair and replacement of damaged vehicle parts.

Galacticar will allow insurance companies to automate the claim adjustment process, dramatically decreasing the amount of time a client has to wait to get an assessment of the damage to their car. It will also help customers decide whether a claim is even necessary, which will hopefully decrease the negative impact an accident can have on their life.

Ayaan Chaudhry, MScAC / **IS:** Alexandros Louizos / **AS:** Professor Richard Zemel and Assistant Professor David Duvenaud

Galaxy AI



Communication Network Analysis for Company Audits

In the world of internal and external auditing, auditors must sort through hundreds, if not thousands of company documents, including financial statements, reports, and more recently messaging systems. These messages between employees and external contacts often contain crucial information on the inner workings of a company, but are often overlooked or missed simply because of the complexity and vast amount of work required to sort through this type of information. In this project we are working with company messaging systems, more specifically emails, in order to build a state-of-the-art analysis tool designed to assist auditors, improve their workflow, and ultimately save companies time and money when going through the auditing process.

This system allows for searching, sorting, parsing and mining of the message content, dates, and users. By performing sentiment analysis, fraud/deceit detection, and clustering based on both the content and network structures of the messages, we automate the process of auditing emails. The framework is designed around extensibility, allowing for integration into virtually any messaging system beyond emails, including third party messaging like Twitter or Slack.

Cole Boudreau, MScAC / **IS:** Brett Kennedy / **AS:** Assistant Professor Frank Rudzicz

CaseWare International Inc.

Designing a Legal Question Answering System

ROSS aims to provide artificial intelligence that can augment the research capacity of the average lawyer and in doing so, increase the time they can spend actually defending their clients. Our focus while at ROSS was to design deep-learning models that leverage state-of-the-art language comprehension techniques, such as word embeddings and Recurrent Neural Networks (RNNs), to provide curated answers to natural language questions. Our time with the company consisted of a cage-match of proposed architectures, each undergoing incremental evolutions until we succeeded in matching our objective. While we primarily focus on a subset of the legal field, our work is undoubtedly generalizable across its other domains. The end result has been a system that outperforms all other competitors in the domain in which it was trained.

MScAC students Colt Chapin, Hao-Wei Lee and Rafael Pedrosa Lacerda de Melo / **IS:** Jimoh Ovbiagele / **AS:** Assistant Professor Frank Rudzicz

ROSS Intelligence Inc.

Distributed Ledger in the World of Digital Identity

SecureKey is designing a digital identity network, based on a distributed ledger platform (Hyperledger Fabric). The solution allows users to transfer their personally identifiable information (PII) from the recognized trusted parties (banks, governments, telcos and credit bureaus) from their own mobile devices. It puts the user back in the middle of the PII transaction, letting them take control over their PII. Hence, the user decides whether to share their PII, and with whom, by giving their explicit consent for the transaction.

By relying on the smart contracts (Chain Code) and powerful cryptography practices, the digital identity ecosystem ensures that any sharable digital identity information is kept private and highly secured during the transaction.

Moreover, the architecture of the identity management service enables triple-blinded transactions. In this case, only the user knows from where and to whom they transferred their identity information. SecureKey's approach offers a completely new method for building a trusted decentralized digital identity ecosystem, enabling the basis of a nation-scale digital identity information market where the user, identity consumer and provider benefit from each and any identity information transaction.

Salavat Nabiev, MScAC / **IS:** Dmitry Barinov / **AS:** Assistant Professor Benjamin Rossman

Trading Strategy Development

Electronica AI is a data science firm, specializing in consulting and trading. The project involves building out a pipeline for training automated models to trade in fiat or cryptocurrency markets. This involves building tools to automatically gather and preprocess up-to-date data (which can be exotic, like sentiment scores). We specify algorithmic trading strategies, which operate on exchanges with the live or historical data. To test the strategies, we build synthetic exchanges which run off of historical data. We use sophisticated black-box optimization techniques (like Bayesian optimization) to tune each of the strategies parameters on the synthetic exchanges. The tuned strategies are validated on held-out data, with methods like walk-forward analysis, and executed on live data with paper trading and money.

Jonathan Lorraine, MScAC / **IS:** Aristotle Andrulakis / **AS:** Assistant Professor David Duvenaud

Extensibility Data Prefetching Strategy for Gmail

Gmail is an email service and a set of email clients developed at Google Inc. To enable smoother user experience for third-party applications which use Gmail as a channel to accomplish tasks, extensibility features are being added to Gmail. Since the extensibility data is fetched on request, there's a delay from the user opening the email to the user being able to use the extensibility features. Some prefetching logic needs to be added to decrease Gmail on the web extensibility data request latency. Some constraints for prefetching in Gmail on the web are: prefetching should not block the main thread; the number of concurrent prefetching requests should be limited; the cache should not add too much burden to heap memory allocation of the entire Gmail web application.

There are two sub-problems. One is to decide the cache size for prefetched result storage. The approach taken was to: identify the limitation of the number of emails a user can read before fresh extensibility data expires. The other is to decide what to prefetch. The approach was to: from past email-reading logs, extract key attributes/features of each email and apply the k-nearest neighbors algorithm to predict the characteristics of emails which will be read, and are thus worth prefetching. After evaluation and adjustment: for the first sub-problem, it was decided to go with a cache size of fifty entry slots; for the second sub-problem, applying the chosen email characteristics gave the prefetching hitting rate of 3.46%.

Yizheng Ding, MScAC / **IS:** Jing Zhang / **AS:** Assistant Professor Roger Grosse

Extracting Supplier Information from the Web

According to the Netcraft Web Survey, there were more than 1.8 billion web pages on the World Wide Web, as of April 2017. The unstructured nature of text content hosted on websites presents various opportunities to apply information extraction techniques, for both commercial and non-commercial use. Our purpose in this research project is 1) To develop web crawlers that will extract information from the world's suppliers' website; 2) Engineer data flow pipelines for fast and efficient extraction of the data for use in machine learning; and 3) To apply cutting-edge machine learning techniques to leverage this information. The extracted information has enabled key features of a supplier discovery and recommendation system. Our effort combines various techniques, including word embeddings, paragraph vectors, approximate nearest neighbour search algorithms, named entity recognition, and neural networks to gather, deduplicate, validate and analyze supplier data. We have produced a database of almost 800,000 suppliers, located primarily in North America, and further applied machine learning techniques to recommend nearby companies offering similar products and services.

Mete Kemertas, MScAC / **IS:** Geoffrey Peddle / **AS:** Assistant Professor Frank Rudzicz

Improving E-Commerce Search: A Data Driven Approach

Flipp is a retail technology company that is reinventing the way people shop. We work with the largest retailers and manufacturers in the world to help them transform their business and connect them with tens of millions of shoppers, through the next generation digital shopping marketplace.

Flipp

Like most online marketplaces, search at Flipp plays a critical role in how users and retailers consume and distribute their content, respectively. The central challenge in making an effective search engine lies in retrieving the most relevant content given a search query. This information retrieval problem starts off with basic text matching, but in order to deliver the most relevant results, more information about the content and users need to be extracted. This project looks at how a data driven approach can be taken to extract these features, through the use of modern machine learning, and use these features to learn how to objectively improve the relevancy of search.

Muhammed Anwar, MScAC / **IS:** Patrick Smith / **AS:** Professor Anthony Bonner

Improving Voice Activity Detection in an Always-On VoIP Application

Switchboard is an always-on VoIP app for iOS, Mac, Slack, and web browsers designed for remote teams. Users can maintain prolonged voice connections with their colleagues while collaborating on projects, enabling greater team cohesion. To ensure that users are not distracted by ambient noise on the other end, as well as to reduce data consumption over metered cellular networks, the app incorporates voice activity detection (VAD) and transmits audio only when the user is speaking.

Synervoz Communications Inc.

We investigated several VAD methods, including spectral feature thresholds, Gaussian mixture models, and neural networks. We concluded that neural networks provide the best accuracy for detecting voice segments in a continuous audio stream. Therefore, we implemented a neural network trained on our corpus of speech samples (collected from Switchboard users) using Core ML, recently introduced in iOS 11 and macOS High Sierra, alongside a non-neural network implementation for backward compatibility with older versions of iOS and macOS. Furthermore, we have built in a number of adaptive power-saving measures to ensure that the app does not deplete the battery by running more Core ML predictions than needed. We have developed a VAD that is accurate, requires minimal user setup, and runs locally on the user's device. It eliminates the need for sending an entire audio stream to a central server for processing, allowing for spontaneous voice chat over metered cellular networks.


Willie Chang, MScAC / **IS:** Jim Rand / **AS:** Assistant Professor Frank Rudzicz

Intelligent Camera Switching

Nureva Inc.

This research project aims to explore the possibility of covering meeting rooms with multiple inexpensive cameras, and intelligently switching and combining camera feeds, based on the best estimated region of interest. Different factors like position of the person speaking, object and action recognition, gaze/attention detection and contextual understanding of the meeting space contribute towards estimating the region of interest effectively. One of the most challenging aspects of this project is the limited time window available for processing all this data. Since we intend to implement this technology to improve the real-time video conferencing solution, we need to process the feeds as quickly as possible, as a delay of more than 100ms is very noticeable to the end users and the video appears to lag.

Nishant Arora, MScAC / **IS:** David Popovich / **AS:** Assistant Professor Sanja Fidler




Interactive Preference Elicitation for Book Recommendations

Rakuten Kobo

Search and book recommendations are the main avenues for content discovery at Kobo. This project focuses on building an interactive query and response system to extract valuable initial book preference information from Kobo customers, to enhance their search and recommendations experience. Currently, personalized book recommendations at Kobo are provided by a recommendation service that relies on many implicit signals – purchases, previews, clicks on the website, etc. Some of these signals, such as purchases, are extremely valuable and result in a very good recommendation performance.

This project focuses on intelligently obtaining explicit book preference information. Such information is especially valuable for new users who have provided no implicit signals; this scenario is known as the “cold-start” problem in the recommendation literature. To align a new user’s tastes with items from our catalogue, we perform unsupervised learning on the Kobo catalogue based on purchase patterns. The preference elicitation process then prompts the user to select clusters, which contain books that he or she likes. To leverage Kobo’s current collaborative-filtering recommendation system, we identify representatives from each cluster, which form the basis of the recommendations. The system allows the user to continuously provide input of items or groups that they like, hereby refining the recommendations.

Mary Elaine Ramos Malit, MScAC / IS: Darius Braziunas / AS: Assistant Professor Scott Sanner



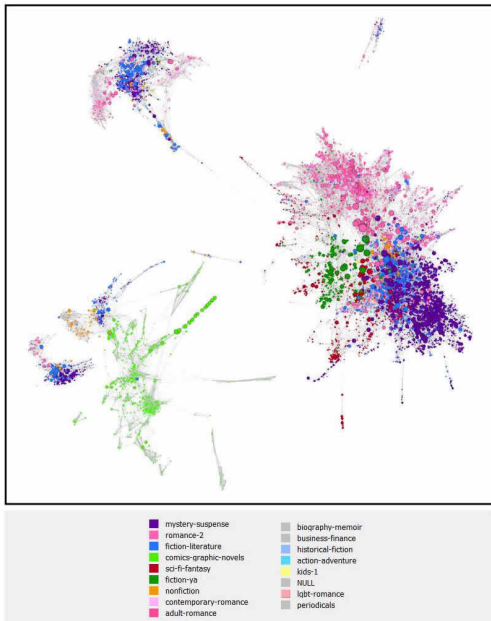
Learning Robotic Grasping for Warehouse Automation

Kindred Systems Inc.

Automated grasping for robotic arms is an open research problem with many applications in healthcare, assistive technologies, manufacturing and logistics. Kindred Systems Inc. is developing products for automatic pick-and-place tasks in distribution centers that provide order fulfillment for e-commerce. One of the tasks that Kindred is trying to solve with efficiency is item sortation. In the sortation task, a robotic arm with an attached gripper is used to pick items from an input location, identify them, and place them in the right output location. This project’s focus is on the first part of this task; picking. It involves finding the best grasping position on the target object, based on sensory information.

The system at Kindred is set up to apply both reinforcement learning and supervised learning techniques to the grasping problem. To adopt a supervised learning approach, a data collection test bed is used to automatically label successful grasps. Leveraging this data, this project involves pursuing different algorithmic approaches and feature designs to achieve higher prediction accuracy. Though a challenging problem, the results show that using point cloud data, 2D features, and simple machine learning models can lead to an effective grasping solution.

Junshu Song, MScAC / IS: Inmar Givoni and Renat Gataullin / AS: Assistant Professor Sanja Fidler



JIT-as-a-Service

The Just-In-Time (JIT) compilers used in modern JVMs improve the performance of Java applications by compiling bytecodes to native machine code at run time. Unfortunately, at the same time they can add significant overhead in terms of processing power and memory footprint. The motivation behind the JIT-as-a-Service project (JITaaS) is to keep all the advantages of JIT compilation while eliminating its problems. The goal is to implement a prototype that can serve as a proof of concept that the JIT Compiler can be effectively run as an independent service, to the point where it could become part of the IBM Cloud service offerings. The prototype consists of two JVMs that communicate with each other over gRPC – a high performance, open source universal remote procedure call framework initially developed at Google. The JVM running the application (here referred to as client) will not do any JIT compilation; instead, it will send all the compilation requests to the JIT compilation service provider (here referred to as server). The server will ask queries about the runtime environment, compile the method, and send a fully compiled body over the wire back to the client. The client will perform necessary relocations to complete the process.

Liberty Application Server was used to benchmark the performance of the prototype. In total, there were 15,678 compilations performed. The most expensive method took 19.6 ms to compile in the JITaaS mode, while it took only 4 ms in the non-JITaaS configuration. The data shows that communication overhead is a major performance challenge that needs to be addressed later.

Fiona Shen, MScAC / **IS:** Marius Pirvu / **AS:** Assistant Professor, Teaching Stream, Mat Zaleski

Machine Learning Meets Economic Complexity

Our Thematic Investing team at CPPIB explores machine learning techniques to further the application of international trade data in constructing systematic investment signals. It builds on the concept of Economic Complexity, spearheaded by Ricardo Hausman and Cesar Hidalgo. It has grown out of development economics through efforts to better understand how and why economies grow, and to identify the key features of those emerging economies that successfully transition up the developed scale.

The social accumulation of productive knowledge is an important factor of economic growth but hard to acquire. Expanding the productive knowledge available in a country is tricky, since industries cannot exist if the requisite productive knowledge is absent. Yet accumulating bits of productive knowledge will make little sense in places where the industries that require it are not present. This “chicken and egg” problem slows down the accumulation of productive knowledge, but creates important path dependencies to predict the economic growth. Countries tend to move into industries that mostly reuse what they already know, since these industries require adding modest amounts of productive knowledge. By gradually adding new knowledge, countries can economize on the “chicken and egg” problem. This is why countries tend to move from the products that they already created to others that are “close by,” in terms of the productive knowledge that they require. Therefore, machine learning techniques help extract the latent representation of countries and products to explain income differences and predict economic growth among the nations of the world.

Yixing Zhang, MScAC / **IS:** Bill Bobey / **AS:** Assistant Professor David Duvenaud

Macroeconomic Data for Financial Statement Prediction

CaseWare has developed a machine learning (ML) pipeline that can predict the likelihood that a particular financial statement will later be modified, or worse, investigated by authorities. But, what does context say about the integrity of a financial statement and the health of a company's finances?

The goal of this project is to extend CaseWare's ML approaches to factor contextual trend data into its predictions. More specifically, the goal of the project is to factor macroeconomic data (such as consumer confidence) and company-specific trend data (such as past performance) into our ML pipeline. Additionally, to better process available signals over time in the trend data and better predict future financial performance, recurrent neural network (RNN)-based approaches were evaluated.

Matthew Brandon Arnold, MScAC / **IS:** Brett Kennedy / **AS:** Assistant Professor Frank Rudzicz



Merchandise Classification with Machine Learning for E-commerce

Personalization of user experience results in a significant increase in sales on e-commerce platforms. Systems that use personalization match up users' attributes with merchandise attributes, such as consumer product categories, target gender and product texture/pattern. Having humans label millions of products that are added daily is expensive and slow. We use NLP techniques and Deep Learning to automatically detect product attributes. Our pipelines are used to facilitate tagging by human taggers, and to tag products without human involvement.

We tag products based on the images and the text provided by the merchants. Product images are classified by product pattern/texture in order to tag a product. We devised a pattern/texture that captures the diversity of products on wish.com, and use a Bilinear CNN architecture to classify images. Our system significantly outperformed human taggers.

Textual descriptions of products are classified by product category and target gender. We extract tf-idf feature vectors, and pass them to an ensemble of classifiers, to output a probability distribution of product category and target gender. We improved performance and reduced runtime by a factor of ten, compared to the baseline.

Van Do, MScAC / IS: Yuli Ye / AS: Michael Guerzhoy, Lecturer and Assistant Professor, Teaching Stream, Matt Medland

Mobile Location-based Insights

Pelmorex (owner of the brands TheWeatherNetwork and MeteoMedia) operates weather information services accessed by all Canadians on desktop computers, mobile apps and television. The mobile app is one of the most downloaded and used apps in the country, within both the Apple/Google ecosystems of smart phones/tablets. Continuous and accurate location data is collected from the apps in a privacy safe way, adhering to Canadian privacy commissioner guidelines and standards such as PIPEDA, with no PII (personally identifiable information) being collected. The location data of millions of users can be utilized by using advanced technologies, such as machine learning and AI, to produce valuable services for consumers and advertisers; personalizing content and advertising based on location history and behaviour. The project explores methodologies and builds new algorithms to both accurately profile the user for their content/advertising preferences, as well as predictive models to help put the right content and ads in front of the user to get better engagement. Better engagement of consumer on both fronts helps Pelmorex to both delight and increase loyalty of its consumers, as well as boost revenue.

Vinay Gera, MScAC / IS: Bala Gopalakrishnan / AS: Associate Professor Mariano P. Consens

Modernizing the Consumption and Management of Standard Response Documents

Standard response documents are summaries of research literature on specific use cases of medications, provided by pharmaceutical companies to assist physicians in prescribing medications safely. Currently, the process for distributing these documents is outdated – the doctor must associate the medication with the respective pharmaceutical company, contact them via email or telephone, and then wait for a static document to be served to them via email attachment or fax.

Doctalk is a tech start-up that is focusing on improving physician processes through a web application platform, beginning with modernizing the centralization and distribution of standard response documents and their updates, while providing valuable tools to the pharmaceutical companies for measuring the impact of these documents within the physician community. As an early stage start-up, it is important that the Minimum Viable Product (MVP) is developed and deployed as quickly and effectively as possible, in order to assert first-mover advantage. To do this, modern requirements engineering is being utilized to ensure the MVP meets the goals and requirements set forth by the stakeholders. The i* framework is being implemented to outline the goals of the system, as well as the MAVO framework to map the uncertainties within the model. As the requirements are unsure and likely to change, the models must be updated regularly as goals and tasks are deemed important or no longer required.

Jamie Beith, MScAC / IS: Dr. Mason Ross / AS: Professor Marsha Chechik

Operating Room Human Traffic Detection

Surgical Safety Technologies (SST) or Blackbox aims to improve patient outcomes by developing data capture technology and processes for the analysis of perioperative factors. This is achieved through a detailed assessment of competence and performance of the surgical team. The assessment includes safety management and error identification within the procedure.

The internship project included the design of a model that would be able to determine the position of the surgical team, and track their movements throughout the procedure to establish the distractions, and possible events caused, during the surgery. Deep learning models, such as Fully Convolutional Networks (FCN) and Region Proposal Networks (RPN), were trained and fine-tuned with a data set put together from the available data from a specific hospital. The precision achieved with the model surpasses those trained on widely known datasets.

Juliana De La Vega Fernández, MScAC / IS: Dr. Teodor Grantcharov / AS: Assistant Professor Babak Taati

Optimizing DiaBits' Predictive Models using Machine Learning Techniques

People suffering from diabetes must regularly monitor their blood-glucose levels to prevent either hypoglycemic or hyperglycemic events. A predictive system that can predict and warn users before the onset of such events can be beneficial to patients suffering from diabetes. Recent approaches to glucose prediction rely heavily on inputs from the user (primarily carbohydrate ingestion, insulin injection, sleeping time). The main problem with these kinds of approaches is that the user may not be able to always fill in this information, or might wrongly fill them (e.g. calories instead of carbohydrates ...). That's why, with DiaBits, we aim to offer a smartphone application, which can predict blood-glucose levels one hour ahead of time, while limiting the need for user input (amount of insulin injected, carbohydrates ingested). To achieve this, we are developing software that integrates continuous glucose monitoring (CGM) data with Wearable Fitness data, to monitor and predict a user's future glucose values. Through this project, our software will be able to distinguish between events of high and moderate glucose variability, and further develop a user-specific learning library storing the habits of the user (eating habits, sport duration, sleeping time). We hope that, by the end of this project, we will increase the software's accuracy compared to our previous approach.

Nicolas Brandt, MScAC / **IS:** Amir Hayeri / **AS:** Assistant Professor Jared Simpson

Polymer – An Ethereum Based DApp Platform

Polymer is an Ethereum based DApp platform and Ethereum is a blockchain-based distributed computing platform, featuring smart contracts with Turing completeness. We have built the Ethereum network for sending and validating transactions in the backend. A Node.js API is built for DApp on Polymer for talking to the smart contracts, querying data, sending transactions and signing messages. The first DApp we built and deployed to Polymer is SCoin Wallet. It's a payment system for supporting P2P money transfer and paying for food in the Digital Factory (DF) cafe. We have built smart contracts for storing employees' bio-information to link to their identity and their SCoin account, so that they can scan their hand at the cashier and money will be automatically deducted from their accounts. All transactions are sent to the Ethereum blockchain, which means they are recorded in an append-only manner and they are immutable, persistent and settled immediately. Business logic is implemented in the smart contract for verifying that the account has enough balance for making that purchase or transfer. SCoin Wallet was deployed on mid-July, 2017 and the team has on-boarded around twenty Scotiabank employees for the pilot version. So far, we have seen 25K CAD loaded and 20K spent on the platform, without any manual effort on settlement and reconciliation. On top of Polymer, we have also built a web application, Ether Block Explorer, for decrypting blocks, transactions and smart contracts information in the network, so that we can better monitor what is happening.

Zhihao Liu, MScAC / **IS:** Henry Kesisyanyan / **AS:** Professor Sam Toueg

Performance Optimization of a Genomic Variant Store for Genotype-phenotype Correlation in PhenoTips

Every person's DNA is made up of approximately six billion base pairs. This genetic code is unique to every individual and is highly conserved; only a fraction of a percent of the genome is different between two randomly selected humans. This small difference is responsible for the variation in heritable disease status between patients and is therefore important in the diagnostic process of genetically-based diseases. In patients suffering from a rare inherited disease, having knowledge of the patient's genetic variants is critical in identifying causation and providing a path for treatment.

This project at Gene42 aims to empower clinicians with the ability to utilize genetic variant data as part of a diagnostic workflow. We have built a tool that allows the upload of variant data and interaction, through a tabular interface. To enable a responsive and interactive interface, the underlying data storage and retrieval engine must be able to perform queries on a massive amount of data with minimal latency time. My work evaluates a variety of different existing technologies, in order to identify the ideal back-end for variant storage and retrieval.

Using this variant storage system, clinicians can upload patient data and quickly elucidate potentially harmful variants by searching, sorting and filtering on variant annotations.

Scott Mastromatteo, MScAC / **IS:** Pawel Buczkowicz / **AS:** Professor Eyal de Lara




PolyRNN-RL: Object Instance Annotation with a Human in the Loop

Annotating object instance segmentation masks is a tedious task that requires tremendous effort and time. In this work, we present Polygon-RNN-RL, a state-of-the-art object annotation method, which uses policy gradients to directly optimize the non-differentiable evaluation metric (IoU). At test time, we used a discriminator and a Graph Neural Network with shared features. This allowed us to select the best polygon from the beam and mitigate the resolution problem.

We showed that our approach outperforms previous methods and baselines by more than 8.5% across all Cityscapes classes, and that without any intervention, Polygon-RNN-RL surpasses the human-level performance for cars on the Cityscapes dataset. To additionally demonstrate the generalization capabilities of our model, we showed several cross-domain experiments on unseen datasets such as KITTI, Rooftop Aerial, ADE20K, as well as on medical imagery. We further introduced an active learning setting, where the model learns from the annotator's corrections, to show significant reduction in annotation time.

David Acuña, MScAC / IS, AS: Assistant Professor Sanja Fidler



Scene Recognition/Parsing

At Layer6 AI we develop state of the art recommendation systems, which incorporate all kinds of data (image, videos, text etc.) to provide relevant and effective recommendations. This project at Layer6 AI involved the development of an open dictionary scene parsing network. A scene parsing/scene understanding system takes an image as an input and assigns a unique class to each pixel in the image. Hence it is able to segment out various objects in the image and recognize them. A new novel algorithm was developed that learned to break down an image into compressed latent representations, through a deep neural network. The algorithm then learned to successfully reconstruct the segmented/parsed image. We plan to integrate this into our recommender systems. The recommender system can then incorporate and use this parsed data to learn visual cues about the preferences of users, and provide better product recommendations.

In terms of application, this work has huge potential in other areas as well, like autonomous driving vehicles, video surveillance, and automated segmentation of internal organs in CTA images.

Himanshu Rai, MScAC / IS: Tomi Poutanen / AS: Professor Richard Zemel



Real Time Virtual Beauty Try On

Modiface

ModiFace is the leading provider of augmented reality (AR) technology to the beauty industry. Our main product offerings provide the ability to perform beauty try-on simulations on live video, and to track the face and facial features in precise detail. Our projects involve hair segmentation, face segmentation, and facial features alignment. For both hair and face segmentation, we used a fully convolutional neural network (CNN), augmented from Mobilenet, to predict masks and achieve fairly good results in accuracy and runtime. In order to achieve more accurate prediction, we needed to preprocess the training images, from basic data selection to more complicated ones, by running facial features alignment and fitting a parabola onto the features to filter out the facial parts. The result could then be used directly on the foundation try-on applications and generate other realistic makeup and hair-dying effects. For the facial features alignment, we experimented with both decision trees and convolutional neural networks. Both models gave accurate facial features alignment and can now be used in many products, whilst becoming the fundamental tool for other projects.

MScAC students Cheng Chang, Yuqiong Wei and Zhi Yu / **IS:** Alex Levinshstein / **AS:** Assistant Professor Sanja Fidler

Temporal Framework for Natural Language Processing with Convolutional Networks

Layer6 AI

This project proposes a new model for document embedding. Existing approaches either require complex inference or use recurrent neural networks that are difficult to parallelize. We took a different route and use recent advances in language modeling to develop a convolutional neural network embedding model. This allowed us to train deeper architectures that are fully parallelizable. Stacking layers together increased the receptive field, allowing each successive layer to model increasingly longer range semantic dependencies within the document. Empirically we demonstrated superior results on two publicly available benchmarks. With the benefits of modern hardware such as GPUs, we also showed that our CNN architecture is over 10x faster than uni-directional skip-thought and over 20x faster than the bi-directional version. Furthermore, by comparing the performance of CNN-pad and CNN-pool, we claimed our CNN model was able to process documents with variable lengths and capture very long range semantic relationships within the document. Decent performance on accuracy, speed and flexibility meant that this model could potentially be applied in different areas such as text feature extraction, sentiment analysis and information retrieval.

Chundi Liu, MScAC / **IS:** Maksims Volkovs / **AS:** Assistant Professor Sanja Fidler


Topological Entity Management System

Scotiabank

Detecting transactional patterns in financial data is important for risk management, customer retention and business development purposes. The objective of this project is to develop a framework for the detection of transactional patterns and their evolution over time. Scotiabank has developed and built an ecosystem that contains both Scotiabank proprietary data and data from third parties. This ecosystem contains data that has been transformed into network data structures, and entity reference data. The network data structures consist of nodes that represent individuals/institutions, and edges that represent interactions between the entities.

Network topology examines the arrangement of nodes and edges within a network and the structure of these relationships. Temporal networks incorporate the time component associated with nodes and edges to better understand the consistency and frequency of a relationship. These areas can be combined to quantify node relationships by measuring their reachability and the strength of their structure. The Topological Entity Management System discovers node relationships expressed in topological shapes. Topological shapes can be defined as repetitive patterns (substructures) within graphs. These include ego-net patterns further classified on the basis of topology, as well as directed digons, triangles, and quadrilaterals. The goal of the project was to develop a flexible and scalable framework to monitor and track topological entities over time. The dynamic system can handle time-evolving data, enabling the modelling and analysis of relationships and trends over time.

Sakshi Garg, MScAC / **IS:** Menaka Kiriwattuduwa / **AS:** Professor Peter Marbach



Training Neural Networks in a Virtual-World using Domain Randomization

Training and testing a deep neural network is a time-consuming and expensive task, which typically involves collecting and manually annotating a large amount of data for supervised learning. This is especially problematic when the task demands expert knowledge, requires labels that are difficult to manually specify, or involves images that are hard to capture. We present a system for training real-world object detectors using synthetic images. To handle the variability in real-world data, the system relies upon the technique of domain randomization, in which the parameters of the simulator are randomized in non-realistic ways to force the network to learn the essential features of the images. We explore the importance of these parameters, showing that it is possible to produce a network with compelling performance using only synthetic data. With additional fine-tuning on real data, we show that such a network can yield competitive performance using much less real data than is typically required. Such a result opens up the possibility of using synthetic data for training real-world neural networks without collecting large amounts of hand-annotated real-world data, which remains a bottleneck for many applications. The approach is validated on both 2D bounding box detection of cars, as well as pixelwise segmentation.

David Acuña, MScAC / IS: Gavriel State / AS: Assistant Professor Sanja Fidler

Understanding What You Say at Work: Speech to Text Recognition

It's hard to engage employees today. Helpful, a video message app, provides a completely new approach, aiming to bring teams together. Video is a more authentic communication tool. At helpful, we try to engage teams and create intuitive and effective environment for our clients. The auto-generated transcription is one feature helpful believe that will enhance app usability. Transcription can help users capture the video content in some situation where it is hard to listen the video. Speech-to-text is the main focus of my internship project.

At Helpful, the presence of proper nouns is key factor in building video comprehensibility, like People and Project Names. Our goal is to provide a reasonable, less obstructive transcription result. We created a learning algorithm for transcription. The context-aware pipeline enables us to get feedback based on user behaviour and information, to increase the comprehensibility of language model. Essentially, the context-aware pipeline build contexts that we can use to incrementally improve the transcription accuracy and behaviour of the application. Moreover, in order to optimize the user experience, we perform post-process on the returned text from ASR (Automatic speech recognizer), such as emoji profanity filter, keyword extraction. Since the result from ASR is normalized, which is hard to read, we predict the punctuation and capitalization for better displaying the transcription.

Yuhui You, MScAC / IS: Graham Toppin / AS: Professor Suzanne Stevenson

Using Machine Learning to Optimize a Workflow Management System

Workflow management frameworks support the creation of task dependencies and make efficient use of resources while running those workloads. Typically, these tasks can be long running processes like Hadoop and Spark jobs, dumping data to/from databases, or running machine learning algorithms. Workflow management consists of mapping tasks to suitable resources, and the management of workflow execution in a cloud environment. Scalability of cloud resources allows real-time provisioning of resources to meet task requirements. The goal of this project is to optimize the resource allocation, using machine learning techniques, in a workflow orchestration framework that manage workloads across an elastic cloud environment. Our proposed approach applies a machine learning algorithm to previous runs of similar jobs, to learn properties of resources required to perform a task. When a new process is initiated, the trained classifier can suggest suitable resources to undertake the specified task. Rubikloud's machine learning pipeline is the backbone of their retail data science products. Computing demand forecasting and product recommendation algorithms require running hundreds of heterogeneous tasks, organized in complex dependency graphs. Currently, the default resource management policy for the scheduler is to allocate the maximum available resources, or previously human determined amounts of resources, to a job. This method does not scale well with the thousands of different configurations of jobs. Often this results in over-committing resources, which leads to paying more to cloud service providers. Adding this feature to Rubikloud's machine learning pipeline would improve the efficiency and scalability of the existing machine learning infrastructure.

Nisal Perera, MScAC / IS: Yuri Bakulin / AS: Professor Eyal de Lara

Working with Mislabeled Data and Improvements to Convolutional Networks for Short Text Classification

This project is broken into two parts: improving classifier performance in the presence of mislabeled data, as well as extending this to make a tool identifying mislabeled items, and improvements to convolutional networks for short text classification. The model for working with mislabeled data uses latent variables, t representing the unknown true label for some input x where we only know the corrupted label c . From this we can create two different model pieces: $p(c | t)$ represents the probability of mislabeling the data and $p(t | x)$ represents the classifier for the true labels. The model is then trained on $p(c | x)$ and never directly observing t : $p(c | x) = \sum_t p(c | t) p(t | x)$

By training the model and rearranging the equation we can also get $p(t | c, x)$ which gives higher predictive accuracy for the true labels which can be used to identify entries that have been mislabeled. The improvements to convolutional networks for short text classification are adding a beginning of term token to the embedding of the input string and instead of mapping characters that it won't be trained on to an "unknown" token, map them to its nearest equivalent character that the model is trained on. For instance 'è' could be mapped to the character 'e', both of which has shown significant improvement in the performance of the trained model.

Darrell Aucoin, MScAC / IS: Bryan Russett and Nicolas Joseph / AS: Assistant Professor David Duvenaud, Professor Radu Craiu



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

CONTACT US

mscac@cs.toronto.edu
www.cs.toronto.edu/mscac
+1 416 978-5180