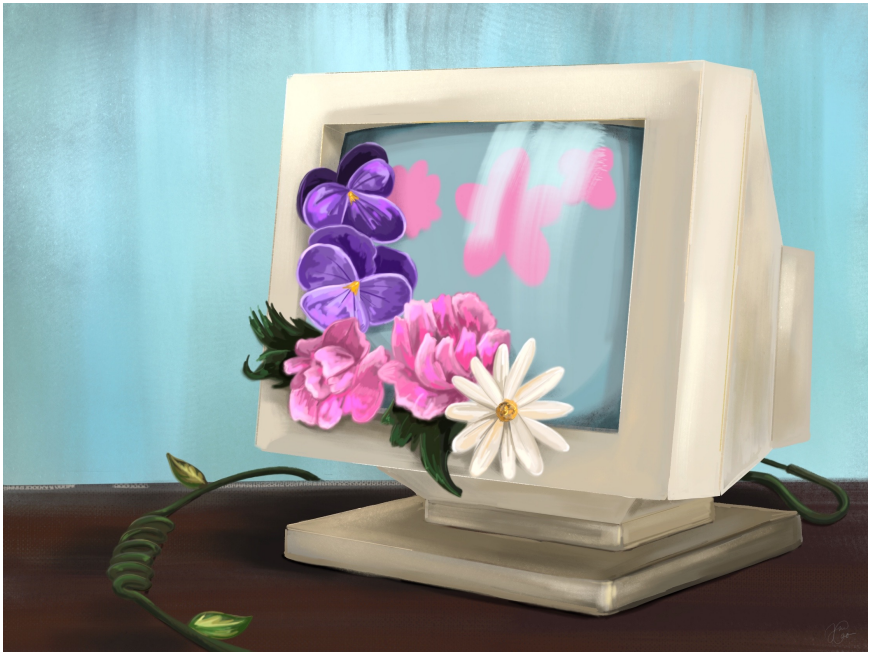

UNIVERSITY OF
WATERLOO

Engineering + EnBus + CS Capstone Design Projects 2024

March 27, 2024

Teams

Accessibility	1	AccessiLooPath
Agriculture	2	Aeroculture
	3	Safi
Housing	4	Bedrock
Healthcare	5	TurnCare
Brain Scans	6	All Joined
Transportation	7	BikeLogs
	8	Composite Damage
Energy	9	GreenMachine
	10	Mercury
	11	Three
Justice	12	Redacted, Cont'd
	13	NewBloom
Rockets	14	TurboPump

Projects

An App for Navigating UW Campus Accessibly
A more sustainable alternative farming solution
Safe milk supply chain in East Africa
A progression dashboard designed to streamline sustainable housing development
TurnCare: A bed attachment for bedridden patient motion to promote bloodflow
BB: EEG To Image Reconstruction
BB: An all-inclusive bike security system.
Damage Detection in Continuous Fiber Reinforced Polymeric Composite Materials
Hydrogen generation from wind turbines
Mercury: The international heavy-industrial instrument tracker
Electricity Maps - open source contributions to help reduce carbon emissions
Improving redaction algorithms
A UI for redacting sensitive documents
Turbopumps for student rockets

1 An App for Navigating UW Campus Accessibly Team AccessiLooPath

Evan Cheer, Jonathan Lanson, Jessica Zhang, Carina Chiu
[Program: MGTE + ECE; TA: Sagar + Akin]

Campus navigation has become increasingly important. Students often find it challenging to navigate through campus with efficient routes without getting lost due to the growth of the campus and the number of buildings, as well as getting the most updated accessibility facilities on campus. Our project aims to develop a mobile app that provides information on the location of accessibility features on campus for individuals with disabilities, including universal washrooms, elevators, and potentially other features such as automatic doors. Our goal aims to enhance the overall campus experience and improve student mobility with the historical and cultural context of accessibility and inclusivity.

The motivation for students to use our mobile app is the desire to navigate campus efficiently and be able to get the most up to date information about the status of accessibility features and washrooms. The app allows students to save time and energy by eliminating unnecessary detours. Our mobile app also aims to cater to different usage modalities by providing visual directions and options to select the accessibility route. This ensures the project meets the diverse needs of the user base.

The user's pain point is the lack of reliable campus navigation solutions. Existing solutions such as WATisRain include information that is outdated and does not satisfy the needs of individuals with disabilities. Additionally, Mappedin's solution at the Student Life Center("SLC") is limited by its physical location; students can only search for routes starting at SLC and must be in front of the kiosk to use it. Our solution aims to provide easy access through students' mobile phones and incorporate accessibility features, such as wheelchair-accessible routes to address specific needs of individuals with disabilities.

As a result, our goal is to develop and deploy the mobile app for the University of Waterloo campus navigation with a focus on accessibility features. User testing and feedback will be carried out for iterative enhancement.

2 A more sustainable alternative farming solution Team Aeroculture

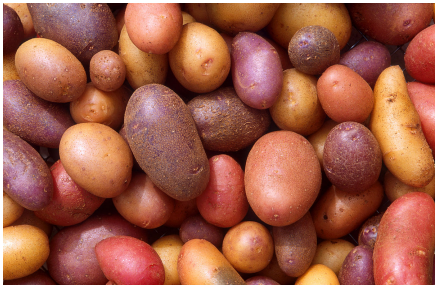

Darby Smyth, Frederick Kwan, Edna To, Leon Han

[Program: Tron + SYDE; TA: Akin + Sagar]

Humans need to be fed, but standard industrial agricultural practices have a negative impact on the health of our planet. Of the 106 million square kilometers classified as habitable land nearly 50% is utilized for agriculture and current agricultural practices account for 70% of freshwater withdrawals. In addition to the vast amounts of land and water that industrial farms require, crops also depend on the use of pesticides and fertilizer to maintain their health which causes chemical runoff that can be detrimental to local ecosystems.

Our solution is an aeroponics system that can be implemented as an alternative to a typical large-scale farm. An aeroponic system is a much more sustainable way to grow plants compared to traditional farming practices as they use less water, do not require pesticides, and can be much more space efficient than large-scale farms. Our solution would be scalable and modular so that a collective system can be implemented in smaller plots of land.

Industrial farming requires a large amount of land and water to grow and maintain crops. The way an aeroponic system operates and is constructed solves both problems. An aeroponic system suspends plants over an enclosed container so their roots grow without the need of a growing medium. The roots of each plant hang in the container where they are exposed to atomizer nozzles that mist a nutrient solution directly onto the roots. The system is also scalable and depending on the physical design, and individual systems can usually be stacked on top of one another to achieve multi-layered farms, reducing the space required to grow many plants.

We will refine goals as we work through the early stages of our project, but some initial ways to guage results we have identified are: - mortality rate of plants in our aeroponics system - cost minimization - through user studies. - a hypothetical large scale implementation

We hope to test this system with people passionate about sustainable and efficient plant growth to observe self-sufficiency, ease of use, and success of growing with little to no prior knowledge of aeroponics systems.

3 Safe milk supply chain in East Africa
Team Safi

Miraal Kabir, Student Two, Student Three, Student Four

[Program: CS + ???; TA: Akin]

Broccoli is a popular and nutritious vegetable that has been cultivated for over 2500 years. While broccoli originated in what is now Italy, it is eaten around the world, with 73% of the modern crop grown in India and China. There are billions of people around the world who enjoy broccoli on a regular basis, creating a huge market opportunity for our BB app.

Blanching is a fast, healthy, and delicious way to cook broccoli, either for immediate eating or as a preparatory step for freezing or sauteing. Broccoli that has been blanched before freezing retains up to 1300% more vitamin C than broccoli that is frozen directly. There are two distinct blanching techniques in common use, and both are supported by BB: boiling and steaming.

A major challenge in blanching is knowing how long to apply the heat. BB uses a feedback control system based on two forms of advanced image processing to ensure perfectly blanched broccoli every time. First, the initial time target is set by measuring the average floret size on the cutting board. Second, during the blanching process, the time target is continuously optimized based on the colour transformation of the vegetable in the pot.

Initial deployment of BB to the public was done during the fall harvest season, in collaboration with the Student Success Office and the St Jacob's Farmer's Market. Farmers reported that broccoli sales to younger adults increased by 57%. Interestingly, the app also drove broccoli sales to newly retired older people, a significant subset of whom are looking for new projects to engage with. App usage was not significant amongst adults of working age. The Student Success Office reported overall improved physical and mental health amongst students who improved their diet by using the BB app.

4 A progression dashboard designed to streamline sustainable housing development
Team Bedrock

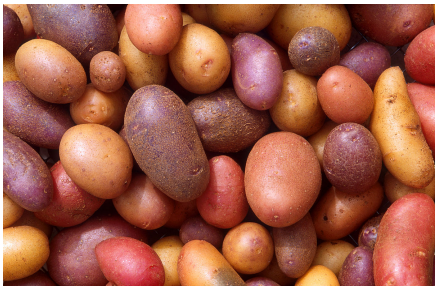

Kefan Cao, Matthew Kanemy, Nathan Leslie$^{\Delta}$
[Program: CS + Econ; TA: Akin]

Bedrock is a progression dashboard designed to streamline the building process for sustainable housing. The dashboard will allow both individual builders and corporations to develop sustainable housing step by step, guiding them towards the best practices in sustainable housing, real-estate development and legal. It will offer tracking abilities throughout each step of the real-estate development process and advise the user towards sustainable development practices. The application will guide the user into thinking about real-estate building in a future-proof way. Building houses this way has a multitude of benefits to offer such as,

- ☐ Energy savings
- ☐ Water savings
- ☐ Improved indoor air quality
- ☐ Reduced maintenance and operation costs

Ultimately, we wish Bedrock would become the standardized pipeline people will use when it comes to real-estate development.

5 TurnCare: A bed attachment for bedridden patient motion to promote bloodflow and fight bed sores. Team TurnCare

Adam Iantorno, Joshua Sewerynek, Noah Coutinho, Lily Tao, Vansh Dhingra, Rishi Vimalendran,
[Program: GENE; TA: Anas]

Pressure ulcers, also known as bed sores, are an injury that breaks down the skin and underlying tissue when an area of skin is placed under pressure [1]. They commonly affect people who have health conditions that prevent them from moving, or cause them to be restricted to a bed or a sitting position. This causes extra pressure to be placed on certain parts of the body, which disrupts blood flow through the skin. Without adequate blood flow, the affected skin becomes starved of oxygen and nutrients, and begins to break down, leading to an ulcer forming.

Despite being seen as a preventable disease, pressure ulcers "affect between 30 and 50% of patients at high and very high risk and susceptibility, especially those hospitalized under critical care". Despite it being known that repositioning a high-risk patient, which means changing the individual's position so that pressure on specific regions is reduced, is an accepted method of pressure ulcer prevention, pressure ulcers still affect 3 million adults in the U.S., costing hospitals and patients an average of $37 800 per hospital stay.

The main method of preventing bed sores in hospitals and nursing homes is having two nurses manually turn the patient onto their side to relieve pressure and allow air flow every couple of hours. However, this method is labour intensive and having two people available to turn those who are bedridden may not be feasible in all cases. For example, if a caregiver does not have the physical strength to turn over their patient, they will need to rely on a different solution.

Currently, a popular solution is an air pump mattress that inflates unevenly throughout to promote airflow and to relieve pressure. However, caregivers are unable to customize any settings on the mattress and patients are still kept in the same position. These air mattresses are also targeted at low to medium

risk patients. For higher risk patients, an alternative solution is required.

Our solution is a bed attachment that can change the patient's position, and an app that allows either the patient or their caregiver to customize the repositioning frequency and method.

The bed attachment will be compatible with patients' mattresses at home, allowing those who do not have professional care to prevent bedsores. The attachment will be able to safely turn the patient on each side temporarily, similar to how nurses turn patients in hospitals.

The attachment will be fully controlled via an app designed for an older user base with accessibility in mind. Patients and caregivers can input the patient's height and weight to determine the force applied. The frequency, duration, and angle of the turning can be customized to each patient's needs. Each patient will be associated with an user profile that can be shared with multiple caregivers to keep track of when the patient has been turned.

The mobile application will have both enterprise (hospital and long-term care home) and direct to consumer use cases that will need to be designed differently. For example, a hospital would have their own account and any nurse would be able to choose a patient profile when dealing with a specific patient, so that the patient's saved settings and history will persist.

6 BB: EEG To Image Reconstruction Team All Joined

Alljoined

Jonathan Xu, Daekun Kim, Charles Liu, Ricky Mao

[Program: CS + SE; TA: Akin]

What is Alljoined?

Alljoined recovers your thoughts using non-invasive brain imaging and deep learning. Electrical activity occurs in different parts of your brain, which can be captured with EEGs. Leveraging transformers and generative models on these scans, we are able to reconstruct what you see and think.

Background and Context

Image reconstruction is the process of taking image stimulus, processing the brain response, and reconstructing the image. Recently, a lot of progress has been made in this field with the rise of diffusion models and transformers. In November 2022, MinD-Vis performed image recostruction from fMRI with a custom train autoencoder + latent difusion model conditioning. In May 2023, MindEye 1 did the same thing but using a two-level architecture with a high-level semantic pipeline and a low-level perceptual pipeline, and used Stable Diffusion as an image encoder. In June 2023, DreamDiffusion adapted the Mind-Vis architecture, but this time with EEG, a more accessible and cost-effective alternative to fMRIs.

Motivation

Inspired by the recent progress in image reconstruction, we wanted to further explore the possibility of reconstructing images from EEGs. Although DreamDiffusion claimed to have reconstructed images from EEGs, they only trained on 40 classes of images. This resulted in a model that acted more like a classifier than a generator. We wanted to see if we could reconstruct images from EEGs with a more diverse dataset, and possibly perform a live demo where we can resconstruct visual imagery from EEGs in real time.

Methodology

The methodology can be broken down into three parts

1. Stimulus & data collection: We collect EEG data from participants while they are shown images. Each session consists of 8 unique blocks, repeated twice. Each block consists of 120 unique images, repeated twice. This results in 920 unique images per session, repeated 4 times in total. Each image is displayed for 300ms with a 300ms interstimulus interval. We also included an additional random 24 back-to-back image repetitions within each block where we ask the participant to press space bar if they notice back-to-back repetitions. We call these trials oddballs, and their purpose is to keep the participant focused so that the EEG data we capture is of higher quality. The EEG data is collected using a 64 channel EEG cap and at 512Hz sampling rate.

2. Preprocessing is a three step process. First we need to filter the data to remove bad trials and oddball stimuli. During this process, we also attach image labels to each of the EEG trials. Second, we need to clean the data. This involves steps like ICA (individual component analysis) where we remove noise caused by actions like eye blink sand jaw clenches. Third, we take our cleaned data and transform it into a format that can be used by the model. This involves steps like downsampling the data to 128Hz and loading the data into a PyTorch dataset.

3. Finally, the model is trained on the preprocessed data. Currently we are using the DreamDiffusion model as it's the easiest one to get up and running while validating that our data is good. DreamDiffusion uses masked signal modeling to pre-train the EEG encoder and fine tunes on a stable diffusion models. By optimizing the EEG encoder and cross-attention heads of the U-Net together, Dream Diffusion enhances the alignment of EEG features with existing text embeddings.

Current Progress

A lot of progress has already been made. We built a $10k high performance cluster (HPC) from scratch and installed a resource manager and task scheduler so that we could open it up to the public and charge users to fund its maintenance. We also conducted weekly paper readings to sync up on the latest research in the field.

This past week (Jan 19, 2024), we rented a house in Pickering and conducted a one week sprint to collect data from the visual reconstruction lab in UofT scarborough campus, preprocess it, and adapt the dream diffusion model

to fine tune on our data. At the end of the week, we presented our findings to Professor Adrian Nestor. You can see our presentation slides here.

Most of our work is pushed to this Github Organization

Future Plans

- ☐ We currently have 8 partcipants to collect data from. Ideally we conduct 4 sessions per participant, where each session is 1 hour long where we can collect EEG/image paired data. So far we only conducted 1 session per participant. We plan to conduct another 2 sessions every Thursday when the UofT lab is available.
- ☐ Preprocessing is currently a manual process. This step is very finicky as we need to extract the EEG data, filter out the noise, and align the EEG data with the image data. Each session takes ~30 minutes. Ideally we automate this process so that we can process more data in less time.
- ☐ On the machine learning side, we'd like to debug the semantic alignment on Brain2Image dataset with DreamDiffusion + MindEye architecture. This part remains somewhat fuzzy as we're considering different architectures that would work well with our limited data but also be able to reconstruct high quality images.

References

References

[1] Chen, Z., Qing, J., Xiang, T., Yue, W. L., & Zhou, J. H. (2022). Seeing Beyond the Brain: Conditional Diffusion Model with Sparse Masked Modeling for Vision Decoding. arXiv. Retrieved from https://arxiv.org/abs/2211.06956

[2] Bai, Y., Wang, C., Xie, S., Dong, C., Yuan, C., & Wang, Z. (2023). Reconstructing the Mind's Eye: fMRI-to-Image with Contrastive Learning and Diffusion Priors. arXiv. Retrieved from https://arxiv.org/abs/2305.18274

[3] Bai, Y. (2023). DreamDiffusion: Generating High-Quality Images from Brain EEG Signals. arXiv. Retrieved from https://arxiv.org/abs/2306.16934

7 BB: An all-inclusive bike security system. Team BikeLogs

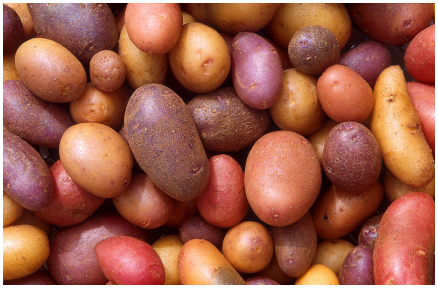

Veronica Leung, Gavin Dan, Jerry Xing, William Ancich

[Program: Tron + SYDE; TA: Sagar]

Bike theft is a prevalent problem in today's society. It is a

particular concern in high traffic areas, such as populous cities and school

campuses. Current big players in the bike security market would include lock

manufactures such as Kryptonite, which debut with their first U-lock in 1971.

It is estimated that bike thefts increase by 400

A large number of bike thefts go unreported to the police, likely due to lack of confidence

in repossession. We aim to be able to provide valuable logged data to track stolen bikes.

A main challenge for victims of bike theft is the tracability. In other

words, there is difficulty in tracking and then proving ownership of a bike.

We are aiming for a two step approach: 1. Preventative measures that deter an attempted

theft to occur in the first place, 2. A tracking system to follow and log the thief's path.

We aim to prevent the theft from occurring in the first place, and even one does occur, to provide

the authorities with enough data to track it down.

Results will be measured in the following stats:

1. Total time and distance tracked
2. Bike thefts attempted
3. Bikes recovered successfully

8 Damage Detection in Continuous Fiber Reinforced Polymeric Composite Materials
Team Composite Damage

Hamad Qazi
[Program: Civil; TA: Anas]

The use of continuous fiber reinforced polymeric composite materials is becoming more prevalent in different manufacturing applications due to its favourable high stiffness and lightweight characteristics. It is particularly useful in aerospace and automotive industries where durable lightweight structures are preferred due to their superior aerodynamic performance.

Having said this, it is important for stress analysts and engineers to be able to predict the behaviour of this material and to determine how the material behaves under fatigue loading. This involves an in depth understanding of crack propagation theory in composite materials.

The crack propagation sequence in composite materials is very different to that of metallic alloys and can not be predicted through the use of traditional Linear Elastic Fracture Mechanics (LEFM). Unlike metallic alloys, composite materials are brittle and will not deform before reaching ultimate failure. This highlights the importance of early detection of micro cracks in the composite material prior to the onset of failure. The focus of this project will be to investigate and further develop the use of a thermographic technique that will be used to monitor the fatigue damage induced within composite materials. The release of heat energy within the composite material matrix upon crack initiation is a key indicator for damage detection. The Infrared Thermography (IRT) technique will involve using an infrared camera to analyze the temperature variation within the composite material as it is being subject to fatigue loading.

The overall aim of this project will be to establish a connection between the thermographic imaging of the material under fatigue load and the resulting failure of the material that will ultimately serve as a pro active damage detection tool.

9 Hydrogen generation from wind turbines
Team GreenMachine

Milind Jain, Jillian Lee, Dafydd Banfield, Ben Hosseinian, Shivam Abhi, Manvir Banwait

[Program: GENE; TA: Sagar]

Renewable energy plants have been used in indusry for many years now. However, due to their nature of generating AC current they cannot easily store energy when generation is higher than supply. One way of storing this energy is to use electrolysis machines to convert the excess electricity into hydrogen which stores the energy inside the hydrogen chemical bonds. Sault Ste. Marie has plenty of renewable energy generation, but lacks the supporting energy transportation/storage infrastructure to fully utilize it. The problem is twofold, with bottlenecks from both off-peak curtailment and saturation of transmission lines.

There exists an untapped potential in Sault Ste. Marie's renewable generation, which could be used to generate hydrogen gas as a means of efficient and green energy storage and distribution. Hydrogen has many uses in the market as it currently stands. It is a great environmentally friendly energy carrier that can be consumed in fuel cells that turn it back into electrical energy when it is needed. Hydrogen is also useful for many industrial scenarios including the steel industry alongside mining and gas locations. There exists a current market demand, as well as a potential future demand for hydrogen as a fuel source or manufacturing aid.

We want a greener future utilizing carbon zero or carbon negative energy sources. Hydrogen is currently an industry leader in green energy potential, with flexibility in its application, making it a safe investment for Sault Ste. Marie. Hydrogen canbe utilized by mining, steel manufacturing, gas, and fuel cells, and is currently seeing projections for increased demand.

Project seems like a significant challenge with multiple facets requiring significant depth and attention to address effectively. This project would not be a simple integration, and financial considerations will be multi-faceted: contractual, infrastructural, technological, startup, training, delivery, and transportation cost to name a few. As well, this option would require comparison with existing and proposed solutions to ensure that hydrogen generation is the most sensible option for the unique application of Sault Ste. Marie.

There are significant challenges with using hydrogen as a fuel source that require technical involvement: Hydrogen generation in a way that is economically favorable, storage, and transportation would all be new challenges to address in Sault Ste. Marie. Additionally, from an engineering perspective, geographic location, climate, curtailment predictability and frequency, maintenance, transportation requirements of equipment, and current technological capabilities would all be factors to consider and address before the project can be viable.

The result will be a thorough engineering business case in favour of or against the generation of hydrogen as a means of energy storage and utilization of curtailed and saturated renewable energy in Sault Ste. Marie. If against, an alternative choice that works best for the region will be justified and presented. In either case, the business case will include a financial breakdown of the project, including projected losses or earnings, potential for fiscal programs, and expected equipment expenses. The results will also include an engineering justification of specific technologies and machines that can be employed to accomplish the goal.

j774lee@uwaterloo.ca

10 Mercury: The international heavy-industrial instrument tracker
Team Mercury

Braden Mah, Eli Vlahos, Kyle Dyck, Nicolas Quintana
[Program: ECE; TA: Sagar]

The demand for sustainable energy products is increasing as climate change and international conflicts threaten energy infrastructure across the globe. In response to this high demand, companies like Tesla are ramping up their energy solutions production. In Q1 2023, Tesla's energy storage deployments increased by 360% YoY with the ramp up of their new Megapack factory in California and they recently announced that they plan on building another Megapack factory in Shanghai later this year (https://ir.tesla.com/#quarterly-disclosure).

With this boom in energy solutions, industrial companies like Tesla will see a large increase in logistics costs in deploying and maintaining their products. Product uptime is extremely valuable for energy solutions and having timely, reliable service to these products is essential.

Servicing industrial energy products safely and reliably often requires heavy-industrial instruments. These instruments can range in size from small handheld devices to massive machines that require multiple shipping containers to transport. When unexpected failures occur, service technicians need these heavy-industrial instruments to quickly resolve issues and avoid downtime penalties. It is imperative that the time between receiving a service request and actually servicing the product is minimized. With instruments being deployed around the globe, they are often far away from logistics hubs where service requests are received. Therefore, knowing the live status of deployed instruments will save employees from spending hours tracking them down thus improving the time-to-service.

Instrument (asset) tracking using manual internal documentation and word-of-mouth is inefficient with a growing fleet of energy solutions deployed on a global scale. Mercury will provide updates on the status of your fleet by attaching a small, Wi-Fi/cellular-connected PCB to each asset, relaying state information to a remote database, and displaying data on an online dashboard. Mercury will also enable remote command execution from the dashboard to

allow users to interact with their assets without needing to physically access them.

Mercury will provide multiple types of information to create an accurate report on the state of the asset being tracked. Information such as location and battery level will be transmitted by default and there will be the option to relay GPIO signals and serial protocol communication messages according to the user's needs. There will also be the ability to read and write data from/to an SD card connected to the asset if applicable.

Lastly, Mercury's online dashboard will allow a user to login with a secure account, view asset information with various filtering options, and subscribe to email notifications on configurable events.

Mercury will give companies control over their globally deployed fleet and they will be empowered to make informed decisions using data collected with our devices. Usage statistics can be used to relocate idle assets to locations with higher demand. Remote information retrieval enables faster diagnostic data collection over long distances. Scalable infrastructure supports a growing fleet of assets. Low-power, battery-supplied devices minimize the frequency of servicing. Lastly, hardware flexibility allows various types of assets to be tracked and data collection configured to suit the company's needs.

11 Electricity Maps - open source contributions to help reduce carbon emissions
Team Three

Wanda Song, Willard Ma, Helen Zhang, Frank Chen [SE]
[Program: SE; TA: Akin]

Electricity Maps (EM) provides vital data that tells the world how it is using electricity. The data is used commercially by companies and other large organizations to help achieve their objectives, while individuals make use of the data for their own personal projects, whether it be for fun or to publish academic papers. We hope as a group to contribute to this existing ecosystem and then offshoot our own research and development so that we can carve a place for ourselves in the fight for a better planet.

Electricity Maps is an open-sourced product that is market-proven and already has many different substantive use cases. The biggest part of it at the moment is the set of API endpoints provided by EM. The main commercial use cases of the product are through the API. The most well known example is that of Google, which uses the geographical electricity data provided by the API to shift its compute workloads to places with greener sources of electricity and meet its sustainability targets. Other companies like Monta (electric vehicle carbon footprint app) and Greenly (carbon emissions accounting app) use the electricity data split by geography as key parts of their product. Companies and individual users alike access the HTTP-based API through an API key obtained through the online key management interface, and then build software programs that make use of the geographical electricity data provided through the API.

EM also has a web and mobile (iOS and Android) app that displays a world map with info on things such as carbon intensity and electricity origin. For individual users, it is a simple and easy to use interface that allows them to browse through some of the data that EM can provide.

The enviornment is a large problem that people and thinkers around the world are constantly struggling with. Having global electricity data neatly provided in one place makes doing research and monitoring the worldwide situation much easier. The open source nature of the product ensures accountability for

the validity of the product iteself. We believe that the data provided by EM can create more opportunities to make advancements in this field.

There are new areas of carbon emissions to discover and with the extension to existing system that we hope to achieve, we want to publish a paper that incorporates the new data that we obtain and presents a novel strategy or strategy in an area where there was previously no data.

Success would be determined by qualitative metrics: For example, being mentioned on client stories would be an amazing result, or having the paper published or presented at a reputable conference.

An acceptable result would be completing a noticeable addition to the open source software and using the data to develop an initial strategy or paper that invites comment and review from the academic community.

12 Improving redaction algorithms Team Redacted, Cont'd

Areena Akhter,
Benn McGregor,[Δ] Kate Granstrom[Δ]
[Program: CS + SE'23; TA: Ahmed]

As of 2022, refugees accounted for 17.23 percent of all permanent resident status holders in Canada. In Canada, the Immigration and Refugee Board (IRB)'s Refugee Protection Division is responsible for adjudicating refugee claims, by adjudicators who are called Board Decision members. The documents detailing these court decisions should ideally be made public, in order to facilitate research into the decision making process behind refugee claims decisions. However, the individuals who claim refugee status in Canada are vulnerable to persecution by their home countries if their information was to be made public, which is why any documents must be thoroughly anonymized before publication.

In collaboration with the Refugee Law Lab (RLL) at Osgoode Hall Law School, I am extending existing research into the development of a large language model (LLM) algorithm that could potentially anonymize unredacted documents in a human-in-the-loop system for automated redaction. The objective of my project is to understand how we would evaluate such a system for its technical and socio-political performance on the redaction task. This investigation will highlight potential assumptions about what the nature of "ground truth" is for a human task, like redaction, as well as introduce new techniques for evaluation of privacy-preserving systems like this algorithm.

The proposed approach to evaluation will be to consider how well our model performs on two competing goals: on one hand, how well it is able to effectively anonymize the refugee's presence in a given document, while also preserving the document's overall utility by not over-redacting it. I evaluate the model against a domain-specific dataset for redacting refugee applications. Then I will apply a set of privacy-preserving evaluation metrics from the Text Anonymization Benchmark to evaluate the model for its performance on two types of sensitive information, personal identification information (PII) and contextual information.

Currently, the IRB uses human annotators for redaction of court documents, a high stakes and time-intensive process that is why few documents end up

released to the public. Current approaches to the anonymization problem tend to be limited to classifying PII, which is a simpler task than understanding how good the model would be at finding contextual, identifying information. This project will prove that this algorithm is a privacy-preserving technique for redaction, and can be judged for success through the publishing of a paper demonstrating our findings.

- ☐ Were we able to establish new evaluation metrics for a redaction algorithm?
- ☐ Did we measure our algorithm against existing approaches, measuring how good they are on privacy-preserving metrics?
- ☐ Did we identify any issues with the ground truth metrics / datasets / training process and improve them?
- ☐ Have we been able to quantify these results in a research paper?
- ☐ Has that paper been peer-reviewed and/or published?

13 A UI for redacting sensitive documents
Team NewBloom

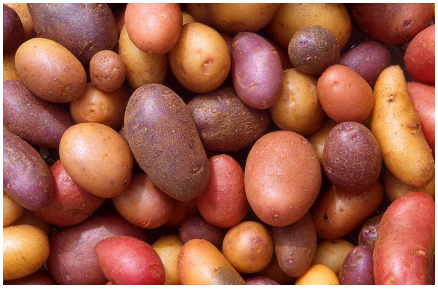

Wendi Yu, Emily Lin, Richard Sun
[Program: SE; TA: Ahmed]

Canada's immigration system provides asylum to politically persecuted refugees. In 2022, the Immigration and Refugee board received 94,246 applications for asylum, and granted 68% of these requests. After the application process, the application's supporting documents are released to the public for the sake of transparency. Because these documents contain personally identifying information (PII) that could endanger these applicants, they must be redacted before release.

The Canadian government employs court redactors who comb through each of these documents and manually redact personally identifying information. This manual redaction is time-consuming and error-prone. Each document goes through initial redaction, and also secondary review to ensure that no information has been missed.

The current process for redacting documents is very manual, and not tailored to the problem space. Court redactors often use everyday text editors (such as Microsoft Word) to redact the documents, and email to send the documents back and forth. Having tailored software that can ergonomically handle redaction and review can greatly reduce redactor burden.

Having tailored redaction software also creates an entry point to leverage a machine learning model to create redaction suggestions, streamlining the redaction process. Machine learning can automate the redaction of not only straightforwards PII, but contextual details as well.

14 Turbopumps for student rockets
Team TurboPump

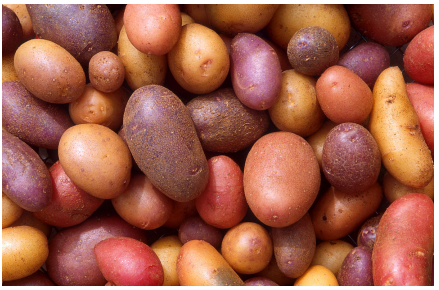

ai2leszk, asotniko, a2hussie, rkobets, l2tomlin

[Program: GENE; TA: Anas]

Turbopumps for use in rocket engines were developed during the second world war, and are a requirement for rockets to leave Earth's atmosphere. As a rocket engine gets more powerful, it requires more propellants. For a rocket engine to operate efficiently, the propellants must be at a high pressure when entering the combustion chamber. To store propellants at high pressure, heavy tanks must be used to store them. By using turbopumps, propellants can be pressurized to required specifications without the need for heavy high pressure storage tanks, resulting in significant weight savings.

Turbopumps are generally considered one of the most complex parts of rocket engine development, and as of yet, have not been used in any student team rockets. By developing this technology, we can pave the way for student design team turbopump development, and open the gates for a new age of Canadian rocket engine development.

The goal of this project is to develop the technology and knowledge at the student team level for future students to use turbopumps on liquid propellant rockets. By creating this technology with our available resources, we are paving a path for future students, and industry partners to follow.

The scope of the project is not to develop a "flight-ready" system, but rather to demonstrate the feasibility of the small-scale pump-fed engine concept by building a working prototype. The goal is that the technology developed, and lessons learned from this project could be used in future work done by the team.

By the end of our project, we intend to have a final test of all components, demonstrating the capability of the technology to operate without significant faults, and meet specified design goals (or as close to them as possible) in a full test using water as the operating fluid. Our goal specifications are to design a turbopump for a 8000lbf liquid rocket engine.

- ☐ Pressure rise: 1000 psi
- ☐ Pump speed: 35000 rpm

- ☐ Fuel mass flow rate: 10.7 lb/s

Cover art by Tina Gao

Student Index

△Kate Granstrom△, 18

a2hussie, 21
Adam Iantorno, 5
ai2leszk, 21
Areena Akhter, 18
asotniko, 21

Ben Hosseinian, 12
Benn McGregor, 18
Braden Mah, 14

Carina Chiu, 1
Charles Liu, 7

Daekun Kim, 7
Dafydd Banfield, 12
Darby Smyth, 2

Edna To, 2
Eli Vlahos, 14
Emily Lin, 20
Evan Cheer, 1

Frank Chen [SE], 16
Frederick Kwan, 2

Gavin Dan, 10

Hamad Qazi, 11
Helen Zhang, 16

Jerry Xing, 10
Jessica Zhang, 1
Jillian Lee, 12
Jonathan Xu, 7
Jonathan Lanson, 1
Joshua Sewerynek, 5

Kefan Cao, 4
Kyle Dyck, 14

l2tomlin, 21
Leon Han, 2
Lily Tao, 5

Manvir Banwait, 12
Matthew Kanemy, 4
Milind Jain, 12
Miraal Kabir, 3

Nathan Leslie△, 4
Nicolas Quintana, 14
Noah Coutinho, 5

Richard Sun, 20
Ricky Mao, 7
Rishi Vimalendran, 5
rkobets, 21

Shivam Abhi, 12
Student Four, 3
Student Three, 3
Student Two, 3

Vansh Dhingra, 5
Veronica Leung, 10

Wanda Song, 16
Wendi Yu, 20
Willard Ma, 16
William Ancich, 10

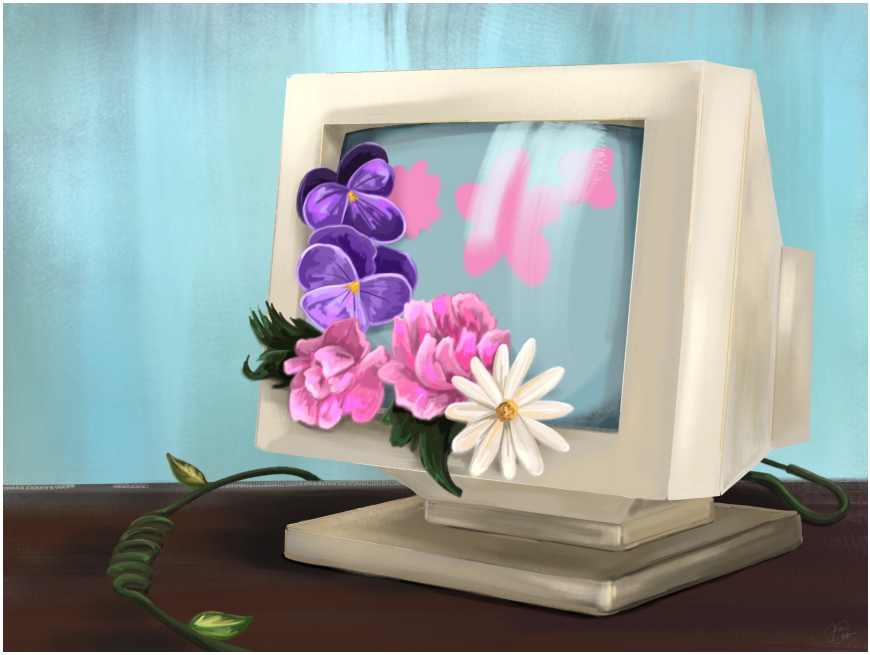

In memory of Grace Chi Hung Leung 1932–2020

UNIVERSITY OF
WATERLOO

FACULTY OF MATHEMATICS
DAVID R. CHERITON SCHOOL OF COMPUTER SCIENCE

FACULTY OF ENGINEERING
DEPARTMENT OF ELECTRICAL AND COMPUTER ENGINEERING