Project ID: 241

SR - Computer Science and Systems Software

Sydney Zhang

EyeScore, A Point of Care Smartphone App for Dry Eye Diagnosis & Prevention

Spontaneous eye blinking can promote tear secretion from lacrimal glands for cleaning, stability and equilibrium of the tear film. Significantly increased eye blinking rate and lid-contact time have been well documented in patients with dry eye disease (DED), the most common eye disease with over 800 million patients worldwide. As a result, blink rate and blink patterns prove be validated biomarkers for DED. In this project, I developed a point of care smartphone app, called "EyeScore". This app allows an iPhone as an imaging and sensing device for 1-minute in-App recording of eyelid movements on iOS platform. The use of eye aspect ratio (EAR) and its various derivatives enables detailed analysis of video frames for determination of eye blinking rate, total lid-contact time and some special DED related eye blink patterns, particularly partial blink. Smartphone videos from ~20 DED patients and normal controls were tested to optimize EAR derivative thresholds for accurate eye blink and partial blink count. Importantly, a clinically relevant algorithm for calculation of "eye healthiness score" was formulated, which takes considerations of eye blink rate, partial blink as well as other demographic factors (such as age, gender, contact lens wearing, etc.). This eye healthiness score can be easily measured at home anytime with a non-invasive manner and identify patients with pre-DED conditions, thus allowing both patients and doctors to monitor eye conditions over time with low-resource setting for early diagnosis and treatment of DED.

Project ID: 242

SR - Computer Science and Systems Software

Frederick Lu

Human Pose Estimation in Determining Karate Pose Accuracy

In martial arts, masters often teach many people at once, and rarely have time to give each student advice. A common example could be a student's arm being at the wrong angle.

We designed a prototype software to help athletes practice forms by themselves. The prototype extracted human pose from camera image and gave immediate feedback based on the difference in joint angles between an attempted image and an ideal image. We chose an open-source neural-network for human-pose-estimation called HRNet. HRNet is lightweight and allowed our prototype to run fast enough on a generic PC without GPU. The "good" pose decision was initially based on percentage error as different actions have different requirements.

As a test, we gathered around 150 images of a black belt and a beginner practicing 3 karate poses. According to our prototype, the black belt achieved good form over 50% of the time while the beginner achieved less than 35% of the time. This showed that our prototype can consistently distinguish good poses from inaccurate ones. We tested using joint angle difference rather than a percentage error as the decision parameter. Using joint angle difference reduced the standard deviation in each data set. Because of this, we modified our software to use joint angle differences.

Our prototype could clearly distinguish a good pose, and could provide feedback at a 0.1fps. Future steps include to automatically detect which pose the user is practicing, and to increase execution speed of the software.

Project ID: 243

SR - Computer Science and Systems Software

Melinda Chang

Albio.js: A Compiler-Powered Reactivity Library for Building Performant User Interfaces

Many modern web programmers utilize declarative user interface (UI) libraries to imbue JavaScript with supplementary infrastructure for interactive development. However, these often interpose extraneous operations at runtime between a user's input and the page's output. Existing solutions that circumvent this behavior are unsuitable for smaller programs because they enforce cumbersome structures that deviate radically from conventional JavaScript.

Albio is a reactivity library engineered explicitly to accommodate small web projects. The initial prototype consisted of a rudimentary compiler capable of analyzing a program and generating JavaScript modules with handlers for primitive reactivity injected into the source code. It computed reactive code during pre-compilation rather than runtime and made granular modifications to the page, improving substantially upon other libraries' overhead. It intentionally lacked components and other constructs that convolute development to maintain ease of use.

I added support for complex data binding, reactive statements, and logic blocks, and I wrote unit tests to rectify functional failings within the codebase. I also designed a set of benchmarks to quantitatively assess its execution times for various standardized operations and modified the compiler accordingly to make optimizations to the rendered code (bitmasking, for instance, yielded a 762% speedup).

Albio is a high-utility alternative to popular libraries for small-scale prototyping and development. It currently demonstrates 142-678% speedups in change propagation as compared to Svelte, React, and Vue. Future iterations of Albio may be further optimized by building a more sophisticated update scheduler, increasing unit test coverage, and expanding the userbase to invite additional real-world feedback.

Project ID: 244

SR - Computer Science and Systems Software

Sophia Guan

A Machine Learning Approach to Understanding the Determining Factors of the Gender Wage Gap

Gender inequality is a complex subject consisting of a variety of issues and nuances. In this project, I chose to study gender income inequality, a prevalent issue in current society. Among the many factors that play a role in the gender wage gap, I focused on the effects of marital status, race, geographical location (by state), age, and years of education. By using these variables to create a model able to predict the hourly wage gap between a woman and their equivalent male counterpart, I could analyze the impact of each variable to better understand the role they play in the income gap. Utilizing income data from the Current Population Survey, I trained and tested five models: a Linear Regression, Decision Tree Regressor, Random Forest Regressor, KNeighbors Regressor, and MLP Regressor. My Linear Regression model found that being a never married worker contributed to a -0.004773 normalized change in wage gap, as well as being a married worker with an absent spouse contributed to a +0.018087 normalized change in wage gap. In general, though, my models had R2 scores below 0.4, so there was virtually no correlation between the variables provided and the predicted hourly wage gap.

Project ID: 245

SR - Computer Science and Systems Software

Gautam Kathiravan

Machine Learning Model for Traffic Behavior Prediction in Intelligent Transport System

Traffic jams and accidents are part of road transportation, especially with manual, semi-autonomous, and fully autonomous cars. These cars use different types of sensors that help sense neighboring cars and lanes and help keep themselves inside the lanes. The current adaptive systems use data from these sensors at a given time. This project focused on predicting the lane-changing behavior of cars (Drivers) using the past state of a vehicle and helping self-driving cars make more accurate decisions while autonomously driving a vehicle.

A recurrent neural network Long short-term memory model was trained to classify the intent of a vehicle as No Lane Change (LK), Left Lane Change (LC), or Right Lane Change (RC) using the time series data with physical, spatial, and social features derived from the Next Generation SIMulation (NGSIM) data. To build such a model, traffic data of vehicles changing lanes with their physical properties played a critical role. Intelligent Transport System with the US Department of Transportation created the NGSIM data on US-101 and I-80 freeway. This data was post-processed to extract features relevant to lane changes which resulted in ~2.8 million rows of data. The training data were sampled randomly to select ~380000 rows with 21 features. As a result, the model performed with a validation accuracy of ~97% and a validation loss of 0.09.

For new vehicle trajectory data, the model accurately predicted the lane change behavior, this can help vehicles make quick decisions in a traffic situation. For the following steps, different RNN models, such as CNN, Transformer networks can be trained with the same dataset and compared for prediction accuracy.

Project ID: 247

SR - Computer Science and Systems Software

Raunak Mondal

Development of Personalized Machine Learning and Data Augmentation Strategies to Diagnose Autism Computer Science and Systems Software

Artificial intelligence (AI) has the potential to revolutionize the way in which autism is diagnosed. One promising application of AI in this context is the use of digital devices to automatically recognize self-stimulatory behaviors related to autism, such as headbanging and arm flapping. However, there are not enough datasets enriched for these behaviors, and the datasets which do exist are very small, making them ineffectual for AI development with neural networks, which are the state-of-the-art approach for machine learning, but which require large datasets to successfully train. To address this issue, it is necessary to explore approaches that are specifically designed to work with smaller datasets. My solution to this problem is to use transfer learning over the I3D neural network, a pre-trained activity recognition model that is trained on the Microsoft Kinetics dataset. Then, I applied data augmentation procedures on the training dataset to correct for overfitting. In addition, I pursued personalized machine learning methods, where the training and test data for the model include the first 80% and last 20% of the same video, respectively. My results indicate that the horizontal flip method is the most performant with an accuracy of 48.78%, while the temporal elastic transformation method is the least accurate method with an accuracy of 24.39%. This project provides a characterization of various data augmentation strategies for activity-based computer vision, thus enabling researchers in the field of AI for behavioral healthcare to more effectively select the appropriate data augmentation strategy when developing clinical-grade models.

Project ID: 249

SR - Computer Science and Systems Software

Shirley Xu

Quantifying Size and Movement of Metal Organic Frameworks Using Computer Vision

This project explores various computer vision techniques to quantify the size and movement of Metal-Organic Framework (MOF) clumps by processing videos of MOFs recorded by a fluorescent microscope. Different thresholding techniques to separate pixels of MOF clumps from the background were experimented with. A two-step static thresholding process is identified and shown to be very effective in identifying pixels of MOF clumps, as the brightness of imaging data captured by the fluorescent microscope varies in the observation time period. Standard OpenCV techniques to find contours were used to identify the boundaries of MOF clumps and calculate their sizes in pixels. The scale bar in the imaging data is identified by its fixed location and unique geometric shape, and the size information is extracted from the scale bar for use in converting the size of MOF clumps in pixels to its real world, physical size. OpenCV's Optical Flow was further experimented on with the original video in the form of extracted contours, convex hulls, and bounding boxes. Optical Flow was also shown to be effective in identifying MOF clump movement in the original video. Further work areas are identified to improve the robustness of the computer vision techniques and calibrate the result from computer vision techniques against other methods of measurement.

Project ID: 250

SR - Computer Science and Systems Software

Benjamin Fan
Edward Qiao

Deep Learning for Solving and Estimating Dynamic Macro-Finance Models

In our project, we use deep learning to perform economic modeling. Specifically, we are given a system of equations describing an economic model, and we use physics-informed neural networks (PINN) to solve these equations for a target function. We design the physics-informed neural network such that its loss function embeds the residuals of the equations to be satisfied; as the loss function is minimized, the residuals of the equations are minimized and the model increases in accuracy.

The PINN takes in an input in the form of a vector, and the input undergoes a series of linear and nonlinear transformations to produce an output which is used to compute the loss function. Furthermore, we use our methodology to simultaneously solve the model and estimate unknown parameters with ease. Performing simultaneous solution and estimation is no more difficult than embedding extra moment conditions into the loss function, and can be done seamlessly.

We apply our technique to two vastly different models, determining the accuracy of a model by measuring the squared error loss between our prediction and a reference solution. Our first model is a model of industrial dynamics with financial frictions, in which we solve for the value function of a bank given equity and productivity. We also estimate the cost of entry and share of labor. Our second model is the standard macroeconomic model with the financial sector described in Brunnermeier and Sannikov (2014), in which we solve for the value of capital and endogenous marginal value of wealth for the banker at an equilibrium. For this problem, we estimate the capital productivity of households.

We find that we consistently achieve an error of less than 1% for target functions and less than 5% for unknown parameters, so we conclude that deep learning is an effective method for economic modeling. This is impactful for several reasons: (1) multi-layer deep learning opens up a new class of highly nonlinear models that can be solved; (2) estimation of parameters and solution of model can be done efficiently in one step; (3) our underlying package automatically applies the state-of-art machine-learning algorithm, so the method keeps improving itself; (4) the method is versatile and can be applied to a vast variety of problems.

Project ID: 251

SR - Computer Science and Systems Software

Collin Giometti

Exploring Interactive Genetic Algorithms with Music Composition

Genetic Algorithms (GAs) are a highly effective method for finding solutions to complex optimization problems. The Interactive Genetic Algorithm (IGA) refers to the incorporation of human evaluation into the optimization process. However, it has been met with several challenges that have led to limited success in its implementation. The objective of this project was to develop an algorithm that could generate optimized solutions around the user's needs, while simultaneously minimizing the number of evaluations. This was represented through musical composition. The initial design was an elementary GA utilizing basic selection and recombination methods. Melodies would be played for the user, who would assign them fitness values based upon the satisfaction of their criteria. During testing, the design exhibited high convergence rates. This was addressed through the expansion of the breeding pool and the usage of stochastic selection techniques, which lowered selection pressure to maintain diversity within the population. The dependence on mutations to introduce diversity was also improved through the utilization of probability distributions, which made them more reliable. The usage of probability distributions in sampling also resulted in a larger, more diverse coverage of the search space. The final design achieved both components outlined in the hypothesis. It discovered an optimal solution on every trial, while the rate of change in fitness also tripled between the first and final design, reducing the number of evaluations needed to reach an optima. Overall, this project improved the effectiveness of the Interactive Genetic Algorithm, allowing more efficient solution generation to meet user needs

Project ID: 252

SR - Computer Science and Systems Software

Agastya Sridharan

Improved Inductive Performance on Small Datasets Using Hyper-Connected Neural Architectures, C. Elegans-Inspired Networks, and Evolutionary Modeling

Neural networks are fundamentally flawed. In even the most trivial tasks like digit recognition, they require tens of thousands of training examples to achieve near-human accuracy. In contrast, toddlers can identify digits with only about a hundred examples.

The generalizability problem has troubling implications for A.I. in the real world. Slight, almost imperceptible alterations to test cases can cause fine-tuned networks to generate absurd conclusions, which is disastrous in critical applications like self-driving and medicine. This project incorporates novel biological structures to increase the inductive performance of neural networks on outliers and small datasets.

We invented a novel class of neural networks termed the MNN, modeled after the visual cortex. The MNN enables recursive feedback mechanisms, in which neurons can communicate with other neurons in any layer. The MNN is trained with a genetic algorithm as a search heuristic that evolves the network's parameters across multiple generations.

As a proof of concept, we test the MNN's generalizability and performance on the MNIST database, with a trainset of just 100 digits. Three networks were incorporated into the MNN: a feedforward DNN, a C. Elegans-inspired mesh network, and a Bernoulli-randomized network. The network achieved 91% test set accuracy with the C. Elegans seed. The best neural configuration outperformed the state-of-the-art Harvard-MIT model on accuracy by 29%, F1 Score by 11%, AUROC by 8%, and AUPRC by 14%. The success of the MNIST testcase is very promising for future small data applications' from classifying social media posts to analyzing genome sequences.

Project ID: 253

SR - Computer Science and Systems Software

Tyler Xiao
Emma Shen

Grow & Give

Student productivity has plummeted following the COVID-19 pandemic. Plunging nationwide test scores signal lower engagement and motivation among elementary to high school students. To address this issue, Grow & Give guides children through planning, setting goals, finishing tasks, and reflecting to accomplish short-term goals. This mobile application includes a special incentive where users accumulate in-game coins from their focused minutes to support real-world nonprofits.

In this project, developers used the Xcode platform to program the backend in Swift and design the frontend with the Storyboard to create a product. A data handling system through the MongoDB Atlas API was used to track user growth and display their progress using SwiftUI charts. To test app effectiveness, students from 2nd to 12th grade participated in a week-long study where the participants were randomly assigned to either the app or the control.

Auto-collected data and a survey show that Grow & Give users saw increased productive engagement and self-impression while the control users saw stagnant results. Feedback from users prompted the developers to improve the aesthetics of the user interface and the intuitivity of interactive elements. Monthly app updates will continue to optimize user experience and widen more opportunities to contribute to social good. Thus, Grow & Give is an app that not only fosters proactive minds ready to contribute to the workplace but also sets a foundation for kids to develop a sense of responsibility to give back to their community.

Project ID: 254

SR - Computer Science and Systems Software

Alexander Proshkin

Investigating the Synergetic Integration of Multifaceted Abstractions with Constraint Programming for Enhanced Time Interval Optimization

This research paper presents a program that addresses the limitations of contemporary commonly used digital calendars by incorporating layers of date and time abstraction onto Google's OR-Tools. Despite the widespread use of digital calendars, with over 70% of adults relying on them as essential tools, these calendars suffer from several limitations, including a rigid association with specific dates and a lack of flexibility. The proposed program uses algorithmic methods and abstraction of integer based constraint programming onto several higher level layers to handle the intricacies associated with interdependent tasks and the frequently encountered constraints imposed on these tasks. The results show that the program offers a simple and efficient solution to the challenges faced by digital calendars, improving their effectiveness and rendering the intended purpose of these calendars achievable.

Project ID: 255

SR - Computer Science and Systems Software

Saanvi Rao

Computer Vision Based iOS App for 2 Way Communication with ASL for the Hearing Impaired

This project aims to address the communication barrier faced by individuals who are deaf or have speech disabilities by utilizing modern Computer Vision technologies. American Sign Language (ASL) is widely used in this community, but not everyone is able to interpret it, leading to a significant communication gap. To bridge this gap, my project is focused on developing a Smartphone App that can recognize ASL and convert it to English speakers in real-time.

The project is composed of two primary components: a Neural Network model and an iOS App. The Neural Network model is designed to recognize ASL hand gestures and was trained using a custom dataset that includes hundreds of pictures of ASL hand signs for each letter in the English alphabet. To ensure the model's robustness, the pictures of ASL hand signs for each letter were captured with slight variations in hand and finger positions. The model was trained using Apple's CreateML tool, which facilitated the training and evaluation process. Once the Neural Network model was successfully trained, it was then integrated into the iOS App that I developed.

The iOS App utilizes advanced Computer Vision technology to detect hand position and landmarks within each camera frame, which is initiated when the Smartphone camera is turned on. The App then leverages the Neural Network model trained to recognize ASL hand gestures, passing the detected hand landmarks as input to the model to predict the corresponding ASL sign. When the App detects a "Space" sign, it combines all the letters detected into a word and uses the speech synthesis capability of iOS to speak the word out loud.

To ensure the accuracy of the Neural Network model, two testing methods were used: an offline test using Apple's CreateML tool and a real-time test using the iOS App. To conduct the offline test, a testing dataset was created by selecting a subset of pictures taken for each ASL sign. The trained model was then validated against this dataset using CreateML, which generated statistics to show the accuracy of the model for each sign in the test dataset. For signs where the model accuracy was below 80%, the model was re-trained with additional images or by increasing the training iterations. The model was also tested in real-time using the iOS App. In test mode, the App displays the predicted letter on the screen, and the results are tabulated to measure accuracy. This real-time testing enables users to see the App's performance in action and provides developers with valuable feedback to further refine the model and App.

One limitation of the current project is that it can only recognize hand poses and not hand actions. ASL utilizes several hand actions to convey words, phrases, and expressions. I plan to continue to enhance the App to be of greater utility to the deaf community and to promote inclusivity and a more accessible society.

Project ID: 256

SR - Computer Science and Systems Software

Rikhil Rao

A Machine Learning Based Pipeline for Automated Single-Cell Quantification of Breast Cancer Marker Evolution from Serial Immunofluorescence Assays

Despite advances in breast cancer diagnoses and treatment, it continues to be the second leading cause of cancer death in women. Recent interest in managing and prognosticating breast cancer has been centered around examining single-cell and spatial expression patterns on molecular imaging, such as immunofluorescence and immunohistochemistry. However, current techniques for this single-cell quantification are disjointed and steps are not formatted to be compatible.

The purpose of this project is to develop a simple, easy to use, end-to-end pipeline to quantify how much a patient's cancer changes at different timepoints in their treatment. To accomplish this, serial cyclic immunofluorescence images acquired over the course of a patient's breast cancer treatment were gathered from the Human Tumor Atlas Network (HTAN). Two images of different samples were chosen from time point 0 and one image from time point 1 was chosen. Each image consisted of 25 unique channels, including DAPI, HER2, and ER. These images were used to train a patient-specific Variational Auto-Encoder (VAE), which was then used to quantify how much the cancer changed. In order to train the model, individual cells were segmented using an open-source tool called MCMICRO, resulting in around 14,000 cells in each image.

By running these analyses, the tumor phenotype of the patient (e.g. HER2 negative, estrogen receptor [ER] positive) was able to be automatically determined. In addition to this the MCMICRO quantification outputs show that the marker expression changed at a later time point as the patient underwent hormone therapy for their breast cancer. For example, the ER expression decreased 2 fold and the Ki67 expression increased 12 fold. Additionally, the VAE latent space plots show significantly different distributions at different time points, implying that the network was able to encode the baseline single-cell expression patterns.

This research shows that automated single-cell quantification of marker expression in breast cancer can be important to track tumor evolution over time, which can be quantified and visualized using a VAE.

Project ID: 257

SR - Computer Science and Systems Software

Arnav Dagar

Where' s Waldo Going - A Bayesian Approach for Localizing a Moving Robot

Autonomous systems have a far-reaching impact: self-driving cars, robots in warehouses, unmanned air vehicles for military use, and rovers for Mars exploration. They use sensors like cameras and LiDARs for perception and localization. LiDARs generate precise maps of surroundings using laser pulses and are often used for Simultaneous Localization and Mapping (SLAM).
Objective: In this project, I implemented and evaluated a particle filter with LiDAR sensor measurements for localizing a moving robot on a custom test station. I also developed a novel camera-based ground truth method for assessing the localization performance.
Method: First, I devised a 3-D LiDAR from a single point TFLuna LiDAR mounted on a 2-axes turntable equipped with slip rings. LEGO motors operated the turntable enabling a 3-D scan. Then I assembled a Raspberry Pi powered modular robot named Waldo(2.0) to mount and operate either a 2-D Slamtech RPLIDAR A1 or the TFLuna based 3-D LiDAR. Blue Dot App remotely controlled the robot with python scripts and acquired scan data (equivalent to 2-D point cloud). I created a 3m x 3m x 30cm test station of custom cut acrylic panels (Waldo's RaceTrack) with rectangular and circular tracks and 30 waypoints using OpenCV generated images printed on vinyl stickers. Next, I devised a novel ground truth method. A Pi camera installed on Waldo observed the ceiling for five numbered and surveyed landmarks. The placements of the landmarks ensured that at least one is in the field of view of the Pi camera from anywhere on the floor. As the robot moves on the tracks, the camera captured the movement of the landmarks in the image. A digit recognition convolutional neural network identified the landmarks in the images. Tracking landmarks in the video frames with respect to a reference frame determined Waldo's location and heading. I recorded LiDAR scans, camera data, color sensor data and motor data for 5 stationary sites, and for 25 runs on the circular track and 25 runs on each of the 4 lines for a total of 130 data sets. I implemented a particle filter in python. Finally, I processed and analyzed offline the acquired data for sensor accuracy, wheel odometry, ground truth from camera and localization using particle filter for different parameters (number of particles, initialization methods).
Results and Conclusions: Though the TFLuna based LiDAR has an inferior accuracy, sampling and scan rate as compared to RPLidar, it still generates good outline of the enclosure. The camera based ground truth method used for A-B comparison showed an accuracy of less than 5cm. The particle filter using LiDAR measurements produced accuracy comparable to the camera method. The RPLIDAR has excellent accuracy and ample sampling and scan rate to extend the project to include more subsystems of an autonomous system (obstacle avoidance, mapping, trajectory generation)

Project ID: 258

SR - Computer Science and Systems Software

Brianna Magtoto

Evaluating Data Augmentation Techniques for Reducing Gender Bias in NLP Using a Pre-trained BERT Model

The research aimed to evaluate the effectiveness of various data augmentation techniques in reducing gender bias in natural language processing models. The study used a pre-trained Bidirectional Encoder Representations from Transformers (BERT) model and compared the results with a baseline model without data augmentation. Six data augmentation techniques were tested, including "gender-neutral substitution", "antonym substitution", "random deletion", "word-sense disambiguation", "context-aware substitution", and "random swap". The hypothesis was that the use of the "gender-neutral substitution" technique would lead to the highest improvement in reducing gender bias in the fine-tuned BERT model. The procedure involved fine-tuning the BERT model with the various data augmentation techniques and evaluating the accuracy and fairness of the models using F1 score and demographic parity. The results showed that the "antonym substitution" technique had the highest accuracy (0.78), while the "context-aware substitution" technique had the highest fairness score (0.60). The conclusion of the study was that while the "gender-neutral substitution" technique did show improvement in reducing gender bias, the "context-aware substitution" technique was found to have the highest fairness score, indicating that it was more effective in treating different groups equally. These findings have important implications for the development of gender-fair NLP models and the use of data augmentation techniques.

Project ID: 259

SR - Computer Science and Systems Software

Nathan Qiu

Metaverse Virtual Classroom for Enhanced Remote Engagement

During the school year of 2020, we were all forced into an unusual computer-based learning environment. A transition to video conferencing software highlighted glaring problems with virtual environments, like lack of face-to-face interaction with peers and proper engagement.

Driven by the lack of existing solutions and the need for a new approach, I developed a 1:1 scale version of a classroom in virtual reality, offering remote users a more immersive and enhanced virtual environment, and offering in-classroom users features only available in VR. My platform would allow both remote and in-person users to engage in face-to-face communication, combining the best qualities of both in-person and virtual environments.

To design my prototype, I first researched and learned the basics of the Unity game engine. I then constructed a virtual scene that incorporated multiplayer and voice chat with PUN Photon Networking, along with user movement. After testing the voice chat, multiplayer, and player movement capabilities, I developed various features such as lighting control, video playback, and virtual whiteboards, which increased the immersive qualities of the virtual environment.

My final virtual environment included all features implemented successfully, creating an immersive VR classroom to facilitate engaging face-to-face communication. There are many ways I can improve and expand upon my current project. One idea is to implement a system that allows users to take a 3D recording of the VR classroom. This is useful for students that have missed a lecture, and need a recording to review the lesson. I also intend to continue my project by more thoroughly field testing my virtual environment, to improve my VR classroom based upon the feedback of the various testers.

Project ID: 260

SR - Computer Science and Systems Software

Woojin Lee

Optimizing the Efficiency and Nutritional Distribution of Food Banks

Despite food banks and pantries offering nutritional aid, 33 million Americans still face food insecurity. This project aims to implement optimization algorithms to improve the efficiency of distributing food in food banks while taking into account individual dietary restrictions. With a dataset aided by information from food banks, different algorithms were compared by computational time and a composite accuracy score. A branch and bound algorithm was found to have the best balance between the two factors. This showcased how food banks can utilize technology to ensure the maximum nutritional output and fulfillment to as many people as possible.

Project ID: 261

SR - Computer Science and Systems Software

Aaron Li

Sober Guardian

The problem of driving under the influence continues to cause an alarming number of car accidents and crashes. To address this, we have developed Sober Guardian, an AI system capable of instantly detecting substance influence in drivers. The system uses facial recognition and visual analysis to compare the driver's characteristics to a pre-trained dataset.

Procedure: Sober Guardian uses facial recognition and visual analysis to determine if a driver is under the influence by comparing their features to a pre-trained dataset. The AI model was trained using a small initial dataset, which underwent preprocessing (photo splitting and organization) and exploratory analysis (recognition of patterns). The training process was carried out using Google Collab, a Convolution Neural Network, and the TensorFlow Lite model, resulting in an accurate prediction of the driver's sobriety status.

Results: We optimized the model's accuracy by adjusting the batch size and epochs during testing (60 epochs, 32 batches, and 0.001 learning rate). The result was a high accuracy rate for detecting drunk drivers (96%) and sober drivers (88%) which exceeded our expectation of 85% overall accuracy. This optimized model was then successfully applied to an external test set.

Conclusion: Sober Guardian has shown that using AI facial recognition and visual analysis can effectively detect and prevent driving under the influence. The ultimate goal is to integrate the technology into vehicles and smart devices to ensure road safety and reduce accidents caused by drunk driving.