Proceedings of the 10th North American International Conference on Industrial Engineering and Operations Management, Orlando, Florida, USA, June 17-19, 2025

Proceedings of the International Conference on Industrial Engineering and Operations Management

Publisher: IEOM Society International, USA DOI: 10.46254/NA10.20250296
Published: June 17, 2025

Road Sense: Utilizing Artificial Intelligence to Detect Drunk Drivers

Aryan Sharma
Student
Queen Elizabeth High School
Calgary, Alberta
Canada aryans0410@gmail.com

Abstract

Drunk driving has remained one of the leading causes of road fatalities worldwide, with an estimated 273,000 deaths annually. Current solutions, such as in-vehicle technologies, outreach activities, and traffic cameras, are either invasive, expensive, or insufficient. In response to this urgent threat, I created *RoadSense* - an AI-powered camera system which uses vehicle behaviour to detect drunk and impaired drivers. It is powered by a Jetson Orin Nano and EMEET 1080p camera, and it includes a 3d printed casing to protect it from environmental circumstances. This system uses 4 behavioural markers to detect impaired drivers following object identification: lane deviation, speed fluctuation, trajectory changes, and proximity violations. Each detection mechanism has been designed using different machine learning models, perspective transformation, and Hough transformation, which allows the system to properly analyze multiple vehicles at once on a road. Testing was conducted in both the real world and simulations in Grand Theft Auto V. This technology has achieved high accuracy in all of its detection mechanisms. In object detection, the model can detect cars, motorcycles, trucks, and buses with an overall accuracy of 73%. In lane deviation, during sunny conditions, the model can predict 95% of lanes. In trajectory fluctuations, the overall F1 score in predicting a drunk driver correctly is 97%. In proximity violations, the model can predict tailgating correctly 95% of the time. At a final material cost of $277, *RoadSense* offers an impactful solution to monitor road safety and reduce human error in enforcement.

Keywords
Deep Learning, Computer Vision, Road Safety, Drunk Driving, Artificial Intelligence

1. Introduction

1.2 The Problem: Drunk Driving

Drunk driving, also called driving under the influence (DUI) or driving while intoxicated (DWI), is a criminal offense that involves operating a vehicle while one's senses are impaired by alcohol or drugs. This impairment affects the driver's reasoning and muscle coordination, affecting their ability to operate the vehicle safely. "Drinking and driving kills 37 people a day in the U.S. — about one person every 39 minutes" according to the National Highway Traffic Safety Administration (NHTSA). Each year more than 11,000 people die in drunk driving accidents in the United States alone.

1.3 What Does Alcohol Do to the Body? How Does it Affect Driving Behavior?

DUI is a crime in which an individual goes behind the wheel after drinking an excessive amount of alcohol. Their consumption amount can be measured by their BAC. This information benefits law enforcement officers and is a leading factor in how the driver behaves. Alcohol impacts various systems in the brain and impairs functions that can be critical for driving. It affects the prefrontal cortex (Judgment & Decision-Making), impairing judgment and impulse control, resulting in misjudging speed, and ignoring traffic rules; the cerebellum (Coordination & Balance) impairing motor skills, making steering, braking, and managing to stay in a lane difficult; the occipital/parietal lobe (Vision Processing), blurring vision, reducing peripheral awareness, and delaying tracking of moving objects, making it difficult to spot road dangers; the basal ganglia (Motor Control and Reflexes), slowing reaction time and disturbs muscle coordination, resulting in uneven braking; the hippocampus (Memory & Awareness), reducing situational awareness and recall, resulting in missed signs and confusion on the road; the amygdala (Impulsivity & Emotions), increasing aggression and overconfidence, resulting in risky decisions and road rage. Studies indicate an exponential rise of the relative risk of a crash with a linear rise of BAC as shown in figure 1. The National Highway Traffic Safety Administration (NHTSA) informs us that blood alcohol concentrations (BAC) in a driver will have the following predictable impact on his or her ability to drive safely.

1.3 Project Objectives

1.3.1 Privacy

Current drunk driving detection systems are extremely invasive and require lots of human interaction. Moreover, since many of these systems involve cloud-based data storage, there is a high chance that user data can be leaked to malicious third-party sources for illegal activities. To address these issues, this project's first objective focuses on making sure that the camera is privacy-preserving. By focusing on external vehicle movements instead of personal driver data, this solution eliminates the need for intrusive in-car monitoring. Additionally, the system processes all data locally on an edge computer to mitigate any further privacy concerns. This approach respects the driver's privacy while still effectively identifying dangerous behaviour.

1.3.2 Cost Effectiveness

Current drunk driving detection methods have not been implemented because of the intense bureaucratic concerns behind them. Many governments, politicians, and vehicle companies have made few efforts to implement existing impaired driving detection technology on a large scale, as it could cost them exorbitant amounts of money to install these devices in every car. To solve this problem, the second objective focuses on making the device cost-effective. This system aims to solve this issue by requiring only one processing unit to monitor multiple vehicles simultaneously. This drastically reduces the cost per vehicle compared to in-car systems (which could cost hundreds of millions), making it far more practical and scalable. By eliminating the need for individual systems in every car, this solution ensures that cities can implement this at a fraction of the cost while still accurately identifying drivers under the influence.

1.3.3 Scalability

Scalability: The impact of drunk driving detection systems is maximized when they can be scaled and deployed in various settings. The 3rd objective is to make sure that the device uses an external camera mounted on a bridge to monitor multiple vehicles' behaviour, rather than individual drivers. To ensure scalability, this device's machine learning algorithm should be generalizable and the technology compact. Finally, the technology used should be protected from weather conditions through a casing.

1.5 Problem Statement

How can I design, program, and construct a scalable, autonomous, and efficient AI-powered device that can externally detect drunk drivers in real-time?

2. Literature Review

2.1 In-Vehicle Solutions

Currently, most drunk driving detection technologies are still under construction. One promising solution lies in the car, where companies are testing breath, facial, and touch-sensing technologies. Figure 2 highlights the efforts of company DADSS, where they are creating breath and touch based technologies.

2.1.1 Breath Based

The systems are in-car integrated breath detectors, which are normally mounted on the steering column. These systems monitor BAC levels via a scan of the driver's breath using a police-style breathalyzer-looking device.

2.1.2 Facial Recognition

The in-car system with an in-car camera scans the driver's face and line of sight by facial recognition with the use of AI. The system searches for the following indicators of impairment: Drooping eyelids or half-shut eyelids (indicative of drunkenness or drowsiness), sluggish blink and sluggish response (since alcohol impacts the nervous system), abnormal head movement (which shows a lack of focus or loss of control).

2.1.3 Touch Sensing
Uses sensors that are installed on the steering wheel, gear shift lever, or ignition button to detect alcohol by skin contact. Uses infrared spectroscopy or an electrochemical sensor to detect BAC by contact with sweat on the skin.

2.1.4 Limitations of In-Vehicle Solutions
One main problem with these detection systems is that they can be highly inaccurate and invasive, as items such as cologne and mouthwash, which contain alcohol, can trigger a false positive warning from the detection system (in breath-based technologies). Coupled with the fact that this driver data can easily be illegally leaked and used for malicious purposes, it hasn't been largely implemented because of the intense bureaucratic concerns involved. Moreover, the high costs of these technologies pose a monetary barrier to implementation. Finally, in some devices, such as car breathalyzers, constant human interaction is required, which can be overwhelming and unnecessary.

2.3 Traffic Cameras in the Status Quo
Currently, traffic cameras are installed around the world. There are 2 main types of these devices: Law Enforcement & Violation Detection cameras detect red-light runners, speeders, and vehicles failing to break at stop signs. Traffic Management cameras are designed to track vehicle density and speed to understand if there are any clogs, detect accidents, and provide environmental condition alerts. These current systems work towards catching simple roadside offenders, however, they fail to analyze driving patterns in vehicles.

3. Methods & Materials
3.1 Jetson Orin Nano
This project was made by programming and training on a personal computer. Downloading videos, sending them to a collab notebook, and getting them processed can take up to 10 minutes, therefore proving inefficient in the long run. If a video is taken of a road, it has to be sent through the cloud, leading to the upload process taking far too long and dangerous drivers under the influence would be out of the frame by the time the video is processed. To solve this issue, an edge computer was sought. These tools are small programmable boards (small "computers") that boast extremely high computing power. They process data closer to the source of the data, rather than in a centralized data center, to reduce latency and improve performance. This project utilized the NVIDIA Jetson Orin Nano. The cost of this device is marketed to be $250 US.

3.2 EMEET 1080p HD USB Camera
The most important component of this project is the camera, which functions as the visual input of the project, providing a high fps playback of the environment. To ensure both reliability and performance, the EMEET 1080p HD USB Camera was chosen. The cost of this camera is $27.

3.3 3D printed case
To enhance the project's durability and reliability, additional measures are to be taken so that it does not get damaged by rain and rough weather. A 3D printed case, based on an online model, has been implemented to serve as an encasing. Additionally, further research is being conducted on waterproofing techniques for electronic components. Various materials are also being tested for added durability and longevity.

4. Detection Methods
4.1 Detect Objects
The very first part of this project, essentially the backbone, is object detection. When a computer processes a video, it cannot tell right away what is what. To help it learn, we can use a technique called deep learning.

4.1.1 Acquiring the Dataset
A computer needs data to learn. The first step in training a deep learning algorithm is finding a dataset to feed to the algorithm. A computer can read this data and extract conclusions using a Convolutional Neural Network. A dataset would be a large file consisting of thousands of images of different types of vehicles (cars, motorcycles, trucks, vans, etc). It is very important to detect all types of vehicles because their type affects their trajectory, speed, acceleration, etc. This algorithm needs to take these factors into account when detecting drunk drivers to avoid any false positives.

4.1.2 Annotating the Dataset

Annotating the dataset is an important step before training the model. Every image/video has to be labelled with information about the objects in it. It defines the spatial extent of each object in an image and draws a simple box around it (called a "bounding box"). Without labelling, the final output can be very inaccurate and messy.

4.1.3 Training the Dataset

After obtaining and labelling the data, we have to train the algorithm. To train on a dataset, it is best practice to use an open-source library to avoid any errors. This project employed the YOLOv8x pre-trained model library. It uses convolutional neural networks to process each image. Convolutional Neural Networks (CNN) are a specialized type of deep learning algorithm that is designed to process all types of visual data, like images and videos. A CNN contains 3 main layers: convolutional layers, pooling layers, and fully connected layers. Figure 3 shows how all of these work together to process and analyze images.

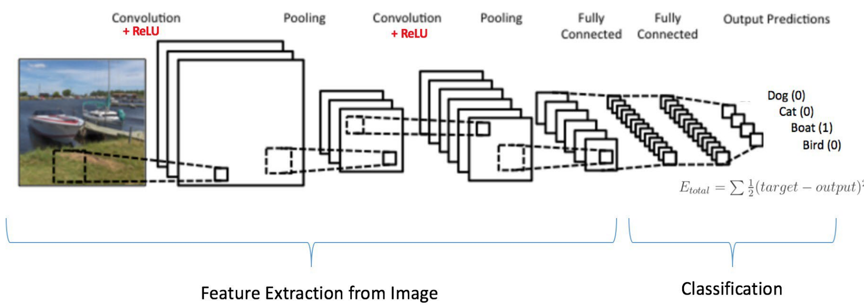

Figure 1. How a Convolutional Neural Network(CNN) processes images

4.2 Detecting Lanes

The next phase of this project focused on lane detection within the selected video footage. When a driver is drunk, they find it tough to identify the "drivable areas" (which are in-between the lanes), which could cause the vehicle to stay in a lane path for a prolonged, unnecessary period. The initial approach to detecting lanes involved using an open source model and tuning it to the project's parameters. To do so, a model that could detect the drivable area in a road was required which would then be used to see if a vehicle is erratic.A comprehensive dataset which fit these requirements was found: BDD100k. Created by the Berkeley AI Research Team, this dataset includes over 100K driving videos, each 40 seconds long, collected from more than 50K rides across locations like New York and the San Francisco Bay Area, in all types of driving conditions. At the outset, this model seemed perfect for the project's needs. However, there were many complications during implementation. The output was not only laggy but inaccurate and noisy, primarily due to the training videos. These clips were shot inside of the car, which means that only if one had a camera (like a dashcam) inside a vehicle would it run smoothly. An external approach, as required by this project, would not work as intended. Figure 4 shows the difference in inferences between internal and external approaches.

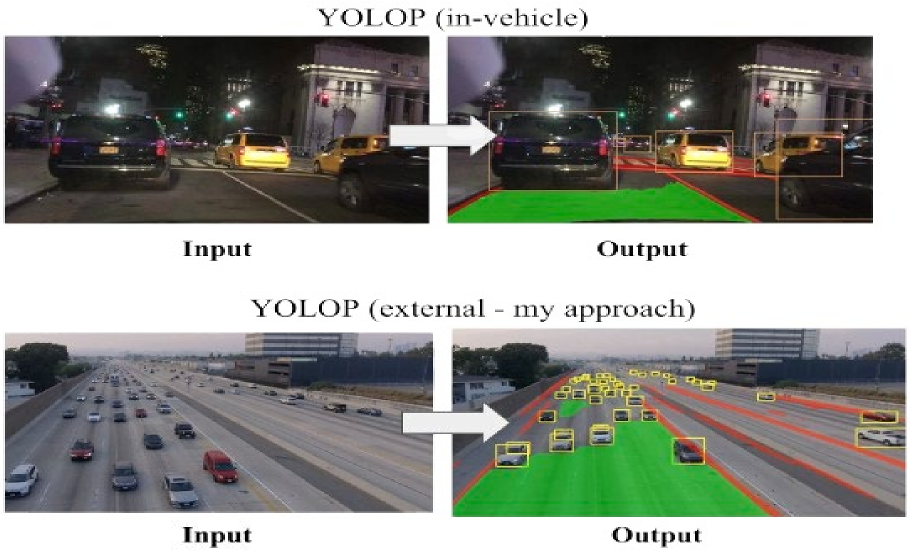

Figure 2. Inference comparison between YOLOP internal and external approaches

Computers can’t pinpoint objects in an image without learning. For lane detection, images are first processed through colour changes and then Hough Transform, a computer vision attribute that detects distinct lines in an image. Detecting lines would be simple if using the “y = mx+b” function, but this equation is not ideal for vertical lines because the slope “m” becomes infinite. Instead, the polar form of the line equation is used:

Where “r” is the perpendicular distance from$r =$ the x *origin cos (toθ)the $+$ line, y *“θ”(sintheta)(θ) is the angle between the normal to the line and the x-axis, and “(x,y)” are the coordinates of an edge point. However, detecting lines in an RGB image is difficult and would produce many false positives (creating lines everywhere). To solve this issue, the program first puts the image, or each frame for a video, into a grayscale. The gray-scale image’s pixels each have an intensity between 0 and 1, compared to the 3 layers between 0 and 255 in a colour image. This not only makes it easier for the Hough transform function but also less computationally intensive on the computer. Next, while it is important to detect as many edges as possible in an image, any possible noise must be filtered out. “Noise” essentially refers to any distractions in an image that can deter the algorithm from the ground truth. Noise can be reduced by a function called Gaussian Blur. It blurs the image so any small pieces of noise would seem like “blobs” so they don't attract too much attention. Next, the algorithm runs each frame through the “canny edge detector”, which finds every edge. Working together, all of these functions are pivotal to making sure that the final Hough transform logic executes completely. Figure 5. Depicts how Hough Transform is performed.

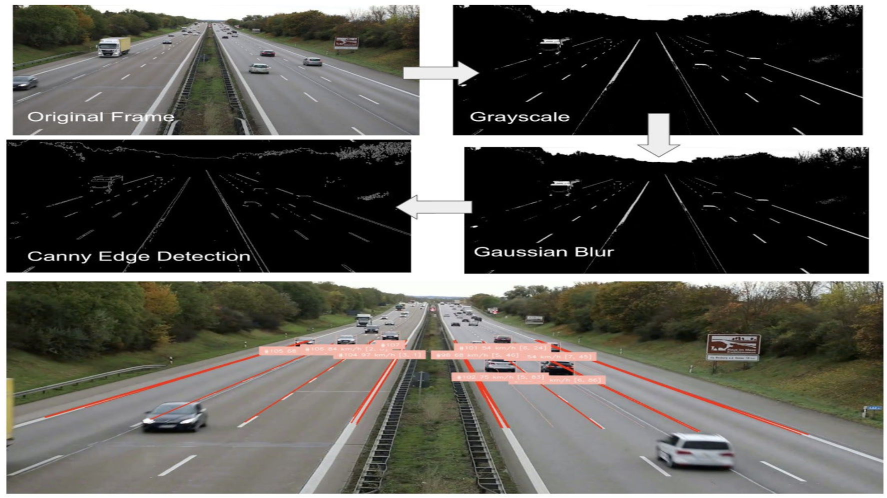

Figure 3. A visual representation of how hough transform is performed

4.3 Detecting Speed

In the status quo, law enforcement uses radar guns to perform speed detection of a moving vehicle. A radar speed gun works by emitting radio waves, which bounce off moving vehicles and return to the gun. Changes in the frequency of the reflected waves(due to the Doppler effect) are analyzed to calculate the vehicle's speed. While this method is generally effective, it can get costly and would make this device bulky and impractical if attached. The first method thought of the solve this issue was to use the speed formula (s = d/t) and the video frames per second. However, after many tries and failures, it was realized that the perspective of the video affects how speed is calculated. The normal equation for calculating speed is distance divided by time, but if we don't have the distance right, then the final calculation will be skewed. The only way to accurately get the distance is to have a bird's-eye view of the road. Initially, a drone was bought to solve this issue. However, there were two main reasons why this idea had to be discontinued: drone regulations in Calgary (where you can fly, how heavy drones need to be, etc) are limiting, and the drone performance itself was not the best – it was wobbly, had a low runtime of 15 minutes, and could pose a safety risk to drivers as it can be visually distracting. However, after many hours of research, a complex Python function called perspective transform was found. It essentially turns an image from a 3D to a 2D plane.

Figure 8 depicts a real-world example of how perspective transform is performed.

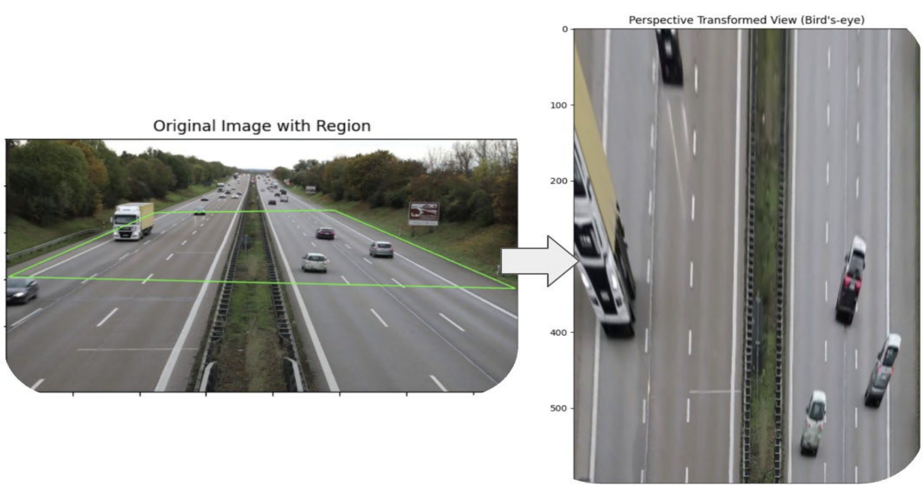

Figure 4. Perspective transformation visualization example

After this, speed was calculated by computing the distance between two video frames and dividing the result by the inverse of the frame rate.

4.4 Detecting Trajectory

A key indicator of drunk driving is the vehicle's trajectory on the road, which the human eye cannot accurately track. AI and computer vision address this by using object tracking libraries such as *BYTETrack*. Following object identification, the bounding boxes on each vehicle are connected using tracklets, a series of frames in which an item has been detected and tracked, via ByteTrack's data association module.

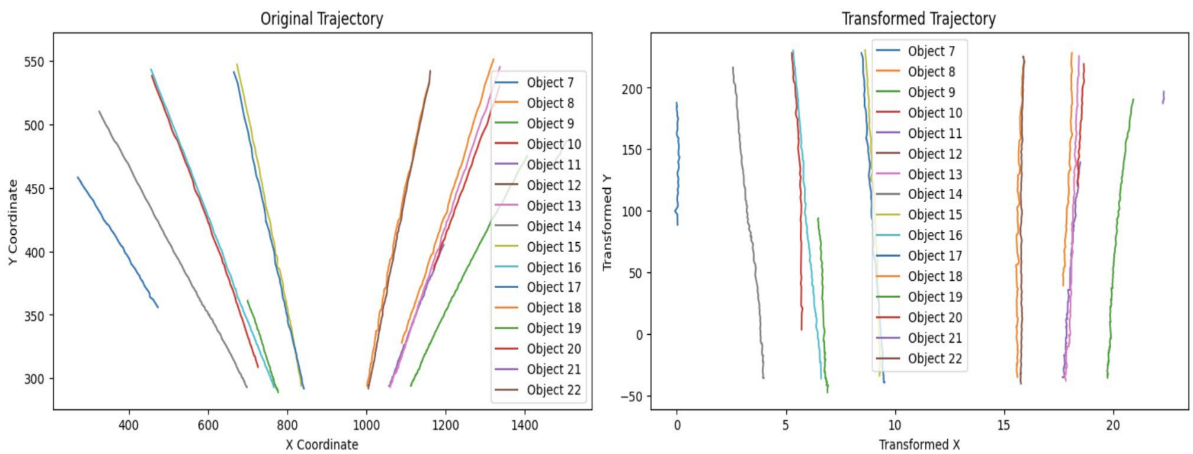

Figure 5. How perspective transform can change trajectories

ByteTrack ensures that objects are precisely tracked over time by linking detection boxes to tracklets. To visualize these trajectories, MatplotLib library was used to plot each vehicle's path. Perspective transform was also applied to align these trajectories with real world coordinates, preventing directions that could affect the model's accuracy. Figure 9 shows how perspective transform changes trajectories to real-world coordinates.

4.5 XY Coordinate Detection

One of the most significant issues with drunk drivers is that they get too close to other cars, roads, or environments; this is called tailgating. To be able to detect these proximity changes, this algorithm uses XY coordinate detection. This works by placing a virtual Cartesian plane on the video input and projecting the location of every car onto proper X and Y coordinates. It can be achieved using homography transformation, which maps real-world points to a top-down perspective, similar to perspective transformation. Each car becomes a point on a Cartesian plane. This enables the system to monitor movement in real-time, detecting dangerous maneuvers like tailgating and hitting infrastructure. Figure 10 highlights the process of how XY coordinate detection works.

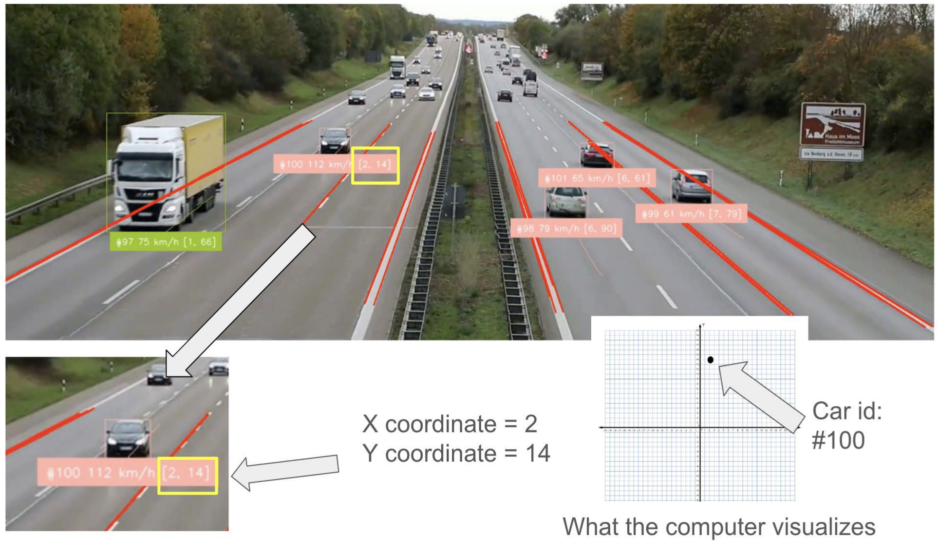

Figure 6. How XY coordinate detection works

5. Results/Findings

5.1.1 Real-World Testing

In the real world, roads can vary from residential streets to the highway. While this device could be adapted to be used on residential roads, prototyping and testing focused mainly on highway conditions for two reasons: there is lots of constant data and since there is a lot of data - there is a higher chance of an intoxicated driver being on the road. To execute this testing strategy, videos of cars driving on the internet were initially used. Once those tests passed, Calgary's highways were surveyed on Google Maps to identify areas where there was an overhead bridge. This way, traffic could be recorded without posing any distractions to the drivers. Once an ideal location was found, the device was set up for approximately 1 hour. For future deployments, collaboration with the City of Calgary to have the device put on and tested on lamp posts will not only ensure safety but also mitigate the issue of theft. Currently, the device can record approximately 50-100 meters of land. Figure 11 shows some examples of the types of roads the device was tested on.

Figure 7. Examples of real-world testing

5.1.2 Simulation Testing

Simulations provide a controlled environment that imitates the real world, allowing for the generation of unlimited, customizable, and accurate data. The chosen simulation needed to fit 3 requirements: it should be able to simulate all types of environments (sunny, dark, rainy, etc), it should have preloaded AI drivers, and it should be multiplayer. After lots of research and testing different simulators, the game GTA V was selected due to its realistic setting, variety in environments, and multiplayer functionality. Using GTA V, scenarios that were recorded in real life were recreated, including standing on a bridge as seen in figure 12. All 4 different detection flags were replicated in nighttime, rainy, and sunny conditions. Finally, a vehicle committing all 4 flags in one clip was replicated.

Figure 8. Simulation testing in GTA V

5.2 Analysis of Each Detection Mechanism

5.2.1 Object Detection

The accuracy of a machine learning algorithm can be seen in Figure 13 through a graph called the Precision-Recall Curve. To understand this, we must understand what precision and recall is first. Precision is the proportion of true positive predictions out of all positive predictions made by the model. It measures how many of the predicted positive cases were correct. Recall is the proportion of true positive cases that the model successfully identified. It measures how well the model captures all actual positive cases. The Precision-Recall Curve visualizes the trade-off between precision and recall at different threshold values. A good model will maintain high precision and recall. As you can see with this curve, the graph is very good and has a high AUC (area under curve).

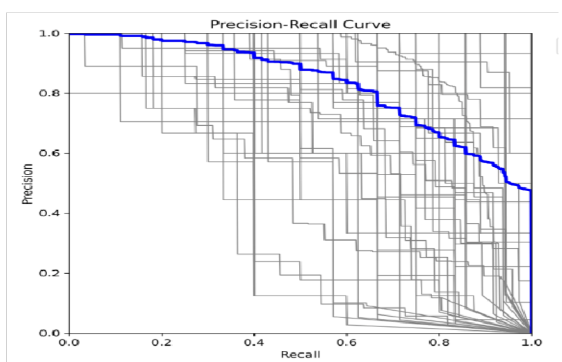

Figure 9. Precision-Recall Curve for the YOLOv8x pretrained model

Table 1. Table showing the Precision, Recall, mAP50, and F1 Scores of Cars, Trucks, Motorcycles, and Buses on the YOLOv8x pretrained model

Vehicle Type	Precision	Recall	mAP50	F1 Score
Car	0.757	0.681	0.752	0.717
Truck	0.687	0.572	0.652	0.624
Motorcycle	0.771	0.717	0.792	0.743
Bus	0.882	0.813	0.886	0.846

5.2.2 Detecting Lanes

When an individual is under the influence, their eye muscle function is restrained. This causes blurry and crossed vision, making it very difficult for drivers to perceive their surroundings, especially at night. It would be tough to identify the “drivable areas” (which is in-between the lanes), and that could cause the vehicle to stay in a lane path for a prolonged, unnecessary period. This behaviour can pose a severe risk to other drivers and pedestrians, and so this algorithm was programmed with the following logic to make sure that this demeanour is flagged. If the car's bounding box is in contact with any one of the hough lines for more than 3 seconds, then 25% is added to the counter. The timer is for 3 seconds or more because a simple lane turn would cause the vehicle to be in contact with the hough lines for approximately 1-2 seconds. One of the core objectives is to make the device implementable and autonomous, implying that the lane detection algorithm has to work for virtually any scenario. Making a model generalizable is important because it is much more convenient to place anywhere, as one does not need to constantly program it to adapt to new conditions. To test the ability of generalization, the Hough Transform algorithm was tested on sunny, rainy, and nighttime driving conditions. If 90% or more of the Hough lines were detected in each video, it was counted as a pass. Otherwise, it would be counted as a “fail”. Below is Figure 14, which details results in this area.

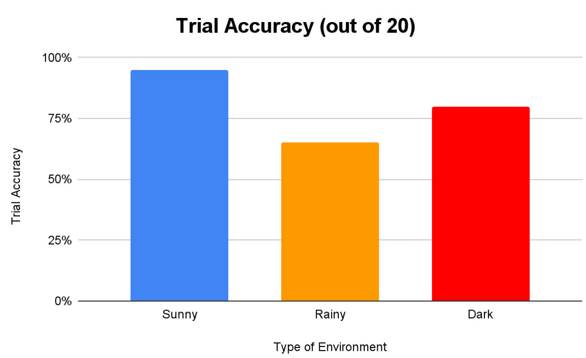

Figure 10. Bar graph representing the effectiveness of lane detection in various environments

5.2.3 Detecting Speed

Drunk drivers can misjudge their vehicle's speed because alcohol slows down cognitive processing, motor skills, and reaction time. This can make the vehicle go over or under the speed limit, or abruptly pick up or slow down speeds. Slower reflexes mean drunk drivers react late to traffic signals, pedestrians, and other vehicles. They may brake too late or too suddenly, causing erratic speed changes. This algorithm detects this by constantly checking on vehicle speed. The vehicle type has to be factored in because vehicles like trucks brake slower than a car for example. The logic in this code states that if a vehicle goes over or under the speed limit for over 2 seconds, then it is flagged. To validate the device’s computer vision based detection, traditional speed measurements were compared with. A real world test involved recording a vehicle driving in the 100-105 km/hr range as indicated by a speedometer and comparing it to the estimated speed recorded by the device from the video footage. This process was repeated three times for accuracy. The result of the speed test are as follows:

Table 2. Results of Speed Testing

Test 1:	Actual Speed: ~103km/h	Detected Speed: ~106 km/h	Difference: ~3km/h
Test 2:	Actual Speed: ~100 km/h	Detected Speed: ~102 km/h	Difference: ~2km/h-
Test 3:	Actual Speed: ~104 km/h	Detected Speed: ~106km/h	Difference: ~2km/h

As shown by the data, the difference is very minimal, which proves that speed detection is very accurate.

5.2.4 Trajectory

Alcohol affects a specific part of the brain called the cerebellum. It is responsible for human balance and movement. When the cerebellum isn't at full capacity, it makes it harder for the driver to steer smoothly, which causes jerky and inconsistent vehicle movements. Moreover, as mentioned before, a drunk driver takes longer to process visual and spatial information. They most likely do not even notice when their vehicle is drifting or swerving in the amount of time it would take for a normal driver to figure this out. After trajectories of cars in the ByteTrack library were detected, which types of trajectories were considered drunk and which were considered sober were then to be deciphered. The first instinct was that a straight line was sober, which is correct, but this assumption failed to realize that cars would still be making left and right turns. When a car takes these types of turns to switch lanes the trajectory isn't straight but slightly curved and ends up forming a cubic function. These trajectories can be seen visually in figure 15. With these three sober trajectories in mind, many images of each were taken and trained in order for the machine-learning model to understand those trajectories as normal. when it comes to a drunk driver's trajectory, it looks like a squiggly line. The machine learning model was trained on Robolflow using the COCO nano model.

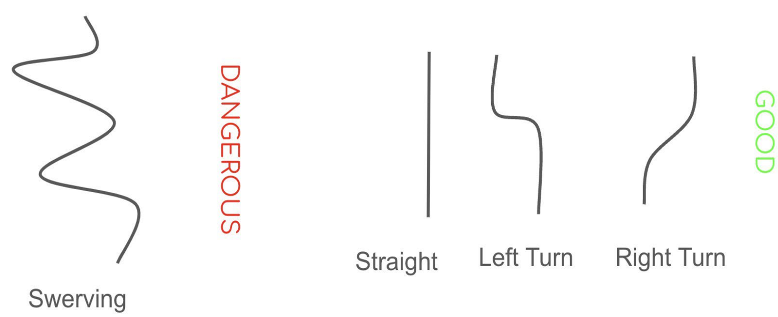

Figure 11. What a drunk vs sober trajectory looks like

Once the machine learning model identifies the trajectory of the vehicle, it adds 25% to the counter if detected as dangerous. The trajectory is only analyzed when the car is just leaving the frame because it works best when analyzing a large period of driving time. Listed are the overall results: mAP = 98.9%, Precision = 96.2%, Recall = 97.2%, F1 Score = 96.7%. Figure 16 depicts the relationship between epochs and precision/recall during training.

Table 3. Results of precision, recall, mAP50, and F1 Score of Straight, turning, and swerving trajectories

Trajectory Type	Precision	Recall	mAP50	F1 Score
Straight (Sober)	0.998	0.997	0.997	0.997
Turn (Sober)	0.989	0.982	0.973	0.985
Swerve (Impaired)	1.00	0.994	0.996	0.997

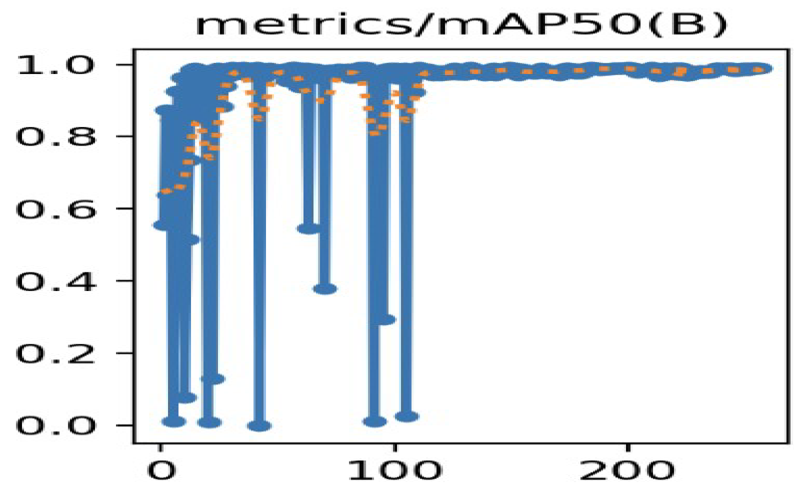

Figure 12. Metrics/mAP50 curve showing the relationship between precision/recall with number of epochs in training

5.2.5 XY Coordinate Detection

Drunk driving disrupts the brain's ability to judge distance. This makes it harder to gauge how close one is to another driver or infrastructure. A drunk driver would think they have more space than they actually do, which could lead to tailgating or near collisions. With the XY detection mechanism, the proximity of each car in relation to the other cars can be detected. The logic behind this relies on calculating the Euclidean distance between two vehicles using their X and Y coordinates. This is done using the formula:

$d = [(x_2 - x_1)^2 + (y_2 - y_1)^2]$

If the calculated distance between two vehicles is below the predefined threshold(50 units), and occurs for over 3 seconds, then the car would be flagged with a 25% added to the counter. To test the effectiveness of the detection flag, the GTA simulation was used by having two cars get very close to each other. This step was to make sure the program flagged this dangerous behavior. This scenario was repeated 20 times. The overall accuracy was 95%, which means that 19 out of 20 trials conducted successfully worked.

5.3 Analysis of Objectives

5.3.1 Privacy

Privacy was one of the most important issues to address in this project. This project successfully made sure that privacy concerns are mitigated by making the device using edge solutions. All of the input data is processed in the Jetson Orin Nano itself, and so 3rd party softwares cannot gain access to the network. Furthermore, the algorithm blurs every car's windshield, making sure nothing in the vehicle is recorded. Currently, the algorithm solely focuses on driving behaviour to successfully detect drunk and impaired drivers.

5.3.2 Cost Effectiveness

Combining the expenses of the Jetson Orin Nano($250), camera module($25), 3D printed casing ($2), Tte final cost of $277 makes this device much cheaper than the possible hundreds of millions of dollars spent for a detection system in every car. It would only need to be placed in select parts of city's, similar to where speed cameras are positioned.

5.3.3 Scalability

This device is fully autonomous and requires no human intervention. In fact, the Jetson Orin Nano is so powerful that it can train itself through a process called unsupervised learning. This means that it can slowly adapt to new data that is getting in the real world and make better predictions without any human training needed. Its compact frame allows it to be implemented on a large scale, in any environment, and the machine learning model behind it makes it "generalizable" or adaptable for any type of situation given to it.

6. Conclusion

6.1 General Conclusion

The purpose of this project was to design and program an autonomous, efficient, and implementable robot that can detect and report drunk drivers in real time. After significant testing and variations throughout each of the 5 detection methods (Speed, XY, Lane, Trajectory, Objects), I have successfully created a device that is $277, fully scalable and performs better than traditional drunk driver detection systems.

6.2 Practical Applications

The practical applications for the drunk driver detection system include law enforcement, traffic safety, and autonomous vehicle technology. This system can be embedded in CCTV cameras or roadside sensors to detect drunk driving and alert the authorities before accidents happen. Police officials can utilize it to select high-risk drivers without using random stops and breath tests, hence facilitating easier implementation of road safety. This system has the potential to save lives, avoid accidents, and bring a revolution in traffic monitoring across the globe.

6.3.1 Device Location

Currently, this device has been tested and implemented on bridges over highways. There is one flaw with this method however - it makes the device prone to theft and tampering. Leaving this system in the open makes it a target for thieves. The next step to solve this issue would be to test the device on high-rise structures such as city lamp posts or intersection areas. This way, not only does it get more road coverage, but it makes the project impossible to tamper with.

6.3.2 Environmental Context Awareness

Erratic driving habits do not always mean intoxication—they may be caused by road conditions, weather, or unexpected obstacles. Upcoming upgrades would incorporate real-time road information, like tight curves, potholes, or heavy rain, to avoid misidentifying innocent movements as drunk driving. Integration of real-time weather API allows the system to adjust its detection parameters based on current conditions.

Acknowledgements

This project would not be possible without the support of the generous grants I received. The Ingenious+ Foundation awarded me $1000, and the Youth Central Foundation awarded me $200 to further my research and testing. This project would not have been financially possible without them.

References

Bajaj, J.S., System and method for driver drowsiness detection using behavioral and sensor-based physiological measures, Sensors, vol. 23, no. 3, p. 1292, 2023. Available: https://doi.org/10.3390/s23031292.

Brooks, A. and Buckey, A, Available: https://www.bankrate.com/insurance/car/drunk-driving/.

What Works: Strategies to Prevent Alcohol-Impaired Driving, Impaired Driving, Available: https://www.cdc.gov/impaired-driving/strategies/index.html.

Dong, M. et al., 'Drinking and driving: A systematic review of the impacts of alcohol consumption on manual and automated driving performance', Journal of safety research, vol. 89, pp. 1–12, 2024. Available: https://doi.org/10.1016/j.jsr.2024.01.006.

The Driver Alcohol Detection System for Safety (DADSS) Program, Available: https://dadss.org.

Perspective Transformation, Available: https://www.sciencedirect.com/topics/computer-science/perspective-transformation.

Drunk Driving, Available: https://www.in.gov/cji/traffic-safety/drunk-driving/.

What are Convolutional Neural Networks?, Available: https://www.ibm.com/think/topics/convolutional-neural-networks.

Evidence synthesis of best practices and effective strategies to reduce drink driving, Available: https://publichealth.jhu.edu/sites/default/files/2024-06/bigrsrapid-responsedrink-drivingfinaljune-2024.pdf.

Keita, Z., An Introduction to Convolutional Neural Networks (CNNs), Available: https://www.datacamp.com/tutorial/introduction-to-convolutional-neural-networks-cnns.

Kingsley, K., et al., In-vehicle technology to prevent drunk driving: Public acceptance required for successful deployment', Transportation research procedia, vol. 72, pp. 2433–2440, 2023.

Kopczyński, J.M.J.K., 'Stand-off detection of alcohol in car cabins, Journal of Applied Remote Sensing, vol. 8, no. 1, 2014.

Kumar, V. et al., Driving to safety: real-time danger spot and drowsiness monitoring system', Soft computing, vol. 25, no. 23, pp. 14479–14497, 2021.

Strategic Plan Overview 2022 – 2024, MADD, Available: https://madd.ca/pages/wp-content/uploads/2023/01/MADD-Canada-Strategic-Plan-Overview-2022-to-2024_FINAL.pdf.

Alberta's Stats on Drinking and Driving, Available: https://maddchapters.ca/parkland/about-us/impaired-driving-statistics/.

Mosel, S., et al., Risks, dangers, and effects of alcohol on the body, Available: https://americanaddictioncenters.org/alcohol/risks-effects-dangers.

NHTSA Issues Advance Notice of Proposed Rulemaking (ANPRM) on Impaired Driving Prevention Technology that will save more than 10,000 lives annually, Available: https://madd.org/press-release/nhtsa-issues-advance-notice-of-proposed-rulemaking-anprm-on-impaired-dri ving-prevention-technology-that-will-save-more-than-10000-lives-annually/.

The HALT Drunk Driving Law – Anti-Drunk Driving Technology for All, Available: https://madd.org/halt/.

Advanced Technology, Available: https://madd.org/advanced-technology/.

The Visual Detection of DWI Motorists, Available: https://www.nhtsa.gov/sites/nhtsa.gov/files/808677.pdf.

Jetson Orin Nano Super Developer Kit, Available: https://www.nvidia.com/en-us/autonomous-machines/embedded-systems/jetson-orin/nano-super-developerkit/.

NVIDIA Jetson Orin Nano Super Developer Kit, Available: https://nvdam.widen.net/s/zkfqjmtds2/jetson-orin-datasheet-nano-developer-kit-3575392-r2.

Raspberry Pi Camera Module 3, Available: https://www.raspberrypi.com/products/camera-module-3/.

Renotte, N., Nicholas Renotte, Available: https://www.youtube.com/@NicholasRenotte.

Roboflow, Roboflow - Real-Time Object Detection., Available: https://www.youtube.com/watch?v=uWP6UjDeZvY.

Comparing DUI statistics of the U.s. vs. Other countries, Available: https://www.sandlawnd.com/dui-statistics-of-us-vs-other-countries/.

Sharma, V., et al., Driver behaviour detection and vehicle rating using multi-UAV coordinated vehicular networks, Journal of computer and system sciences, vol. 86, pp. 3–32, 2017.

How do red light cameras work?, Available: http://progressive.com/lifelanes/red-light-cameras/.

Van Dyke, N. and Fillmore, M.T., Alcohol effects on simulated driving performance and self-perceptions of impairment in DUI offenders, Experimental and clinical psychopharmacology, vol. 22, no. 6, pp. 484–493, 2014.

Drunk driving, Available: https://en.wikipedia.org/w/index.php?title=Drunk_driving&oldid=1274055399.
Zhao, X. et al. Identifying method of drunk driving based on driving behaviour, International journal of computational intelligence systems, vol. 4, no. 3, pp. 361–369, 2011.

Biography

Aryan Sharma is a grade 9 student from Calgary, Alberta, attending Queen Elizabeth High School. He has been creating innovative technologies since grade 5, earning top honours at the Canada-Wide Science Fair (CWSF) in 2023 and 2024. Outside of academics, Aryan loves to spend time with his friends and family, play ice hockey, and hike outdoors. His research interests include Artificial Intelligence, CAD Modelling, and Environmental Protection.