

ICanC: Improving Camera-based Object Detection and Energy Consumption in Low-Illumination Environments

Daniel Ma^{*}, Ren Zhong^{†,‡}, and Weisong Shi[‡]

^{*}Khoury College of Computer Sciences, Northeastern University

[†]Department of Computer Science, Wayne State University

[‡]Department of Computer and Information Sciences, University of Delaware

Abstract—This paper introduces ICanC (*pronounced "I Can See"*), a novel system designed to enhance object detection and optimize energy efficiency in autonomous vehicles (AVs) operating in low-illumination environments. By leveraging the complementary capabilities of LiDAR and camera sensors, ICanC improves detection accuracy under conditions where camera performance typically declines, while significantly reducing unnecessary headlight usage. This approach aligns with the broader objective of promoting sustainable transportation.

ICanC comprises three primary nodes: the *Obstacle Detector*, which processes LiDAR point cloud data to fit bounding boxes onto detected objects and estimate their position, velocity, and orientation; the *Danger Detector*, which evaluates potential threats using the information provided by the Obstacle Detector; and the *Light Controller*, which dynamically activates headlights to enhance camera visibility solely when a threat is detected.

Experiments conducted in physical and simulated environments demonstrate ICanC’s robust performance, even in the presence of significant noise interference. The system consistently achieves high accuracy in camera-based object detection when headlights are engaged, while significantly reducing overall headlight energy consumption. These results position ICanC as a promising advancement in autonomous vehicle research, achieving a balance between energy efficiency and reliable object detection.

Index Terms—object detection, autonomous vehicles, LiDAR, cameras, computer vision, energy conservation, sensor fusion

I. INTRODUCTION

In low-illumination environments, vehicle headlights are essential for ensuring safety and visibility, especially for drivers and passengers. Autonomous vehicles (AVs), which rely heavily on camera sensors for navigation, also need headlights to improve visibility and provide clearer images for decision-making. However, headlights are not always necessary for safe operation, particularly when utilizing versatile sensors such as LiDAR, which can offer consistent performance regardless of lighting conditions. Minimizing headlight usage in AVs presents a promising strategy for reducing energy consumption without compromising safety or detection capabilities.

The most prevalent approach to AV perception typically involves camera sensors, which capture 2D images of the environment. These images provide detailed visual information but lack depth perception and are heavily dependent on lighting

This research was sponsored by the National Science Foundation Research Experience for Undergraduates Program.

	Day	Night	Night w/Flashlight
LiDAR Detection Rate (%)	100	87.5	100
Camera Detection Rate (%)	100	25	100

TABLE I: Percent of pedestrians detected by LiDAR and camera algorithms, as recorded during experimental observations under varying conditions.

conditions. This reliance on cameras can become problematic in low-light or adverse weather conditions. For instance, in 2019, a Tesla vehicle operating in autopilot mode collided with a truck before sunrise, leading to a fatal accident. The incident was attributed to the Autopilot system’s inability to detect the truck in time due to poor visibility in low-light conditions. As Tesla’s system relied solely on cameras, it failed to “consistently detect and track the truck as an object or threat as it crossed the path of the car [1].” This tragic incident underscores the limitations of camera-based perception in low-illumination environments and highlights the need for more robust detection systems in these conditions.

To overcome these limitations, the combination of multiple sensors—specifically cameras and LiDAR—holds significant promise. LiDAR, or light detection and ranging, emits laser beams in a defined field of view, often in a 360-degree cycle, and measures the time it takes for the beams to return. This process generates a 3D point cloud, where each point represents coordinates and additional attributes. Due to its 3D nature, this point cloud contains depth information that camera images lack. Furthermore, these lasers operate independently of the illumination state, allowing the advantage of nighttime usage. This trait is evident from experimental observations, which show that LiDAR detection rates only degrade approximately 12.5% in the dark, while camera detection rates degrade by 75% (see Table I). However, LiDAR has its own limitations as well. Firstly, lasers are similarly prone to noise, or unwanted and irrelevant variations in data, from adverse weather conditions. Additionally, while LiDAR sensor costs have decreased over time, they remain more expensive than camera sensors [2]. Despite these challenges, the fusion of camera and LiDAR data allows for the strengths of each sensor to compensate for the other’s weaknesses, making it an ideal solution for

enhancing object detection in low-light environments.

By fusing the data from both cameras and LiDAR, headlight usage may be optimized while maintaining high detection performance in the dark. When no danger is detected, the headlights can remain off, saving energy. When a potential threat is identified, the headlights can be activated to enhance camera performance. This selective use of headlights reduces unnecessary energy consumption without sacrificing safety or detection reliability.

To implement this concept, this project introduces ICanC—a novel system designed to optimize headlight usage while ensuring effective object detection in low-light environments through sensor fusion. ICanC consists of three primary components: Firstly, the *Obstacle Detector* node, which processes LiDAR point clouds to identify objects and tightly fit 3D bounding boxes around them. This data is passed to the *Danger Detector* node, which analyzes bounding boxes to identify potential threats in each frame. Analysis of data is then passed to the *Light Controller* node, which activates headlights when danger is detected, enabling the camera to capture clearer images of the surroundings. Through sensor fusion of LiDAR and camera data, ICanC ensures that detection performance remains strong in low-light conditions, while headlight usage is minimized to optimize energy consumption. This approach not only improves safety in autonomous driving but also supports the broader goal of sustainable transportation by reducing the environmental impact of electric autonomous vehicles.

Furthermore, this project aims to demonstrate the Vehicle Programming Interface (VPI) and its potential to support the efficiency of programming with autonomous vehicles. The VPI was developed as a response to the growing interdisciplinary complexities of automotive systems, which make developing applications for AVs more challenging [3]. Much of the logic implemented in ICanC can be abstracted into the VPI and employed to solve various problems, thus promoting the goal of accessibility of AV programming.

In conclusion, this paper makes the following contributions:

- 1) Defines the low-illumination environment problem for object detection in AVs.
- 2) Introduces ICanC, a novel system that improves low-illumination object detection and simultaneously reduces energy consumption for AVs. ICanC is open-source ¹.
- 3) Evaluates the system in both extensive real-world and simulation tests.

This paper begins by reviewing related literature in Section II before presenting an in-depth explanation of the ICanC design in Section III. Section IV then discusses the implementation of ICanC in physical and simulated environments. Section V showcases results gathered from extensive experimentation and testing, and Section VI then discusses these results, potential improvements, and future works. Section VII then summarizes and concludes the project. Finally, Section VIII offers acknowledgments to those who provided invaluable support during the project.

¹The code for ICanC is available at https://github.com/danielma4/ICanC_v2.

II. RELATED WORKS

Related works can be classified into two research areas: improving low-illumination object detection and energy conservation in autonomous vehicles.

A. Low-illumination object detection

Current low-illumination AV research utilizes a multitude of techniques to improve performance of detection algorithms. One such approach is *sensor fusion*, or using multiple sensors in conjunction. In particular, synergizing camera images with thermal imaging has shown success in surveillance object detection. Thermal sensors record changes in the temperature of an environment, allowing for an image largely unaffected by the absence of light. By combining images outputted by cameras and thermal imaging, the best of both sensors can be meshed into an improved snapshot [4]. Another sensor fusion technique is gated imaging. Laser-gated imaging operates by emitting laser photons in tandem with a synchronized imaging sensor. This photon offers a brief source of light and illuminates the area. The reflected light returns to the imaging sensor, offering information about the surroundings. Research has experimented with gated imaging in autonomous vehicles, providing advantages of 2D and 3D object detection and performing well in both low-illumination and adverse weather conditions [5]. While ICanC builds on sensor fusion delineated in the above research, this system explores the potential of LiDAR rather than thermal or gated imaging. Furthermore, the ICanC system employs more straightforward logic by keeping the imaging of both sensors separate rather than combining the two into a fused image or sensor. This distinction allows ICanC to leverage the unique strengths of each sensor independently, offering a simpler yet effective approach to enhancing object detection in low-illumination environments.

Another approach is *image preprocessing*. This technique can be accomplished through numerous techniques, including artificially brightening images and extracting pertinent information. Research utilizing Gaussian Process regression has shown success in retrieving relevant information rather than aesthetic restoration; however, artificially altering images inevitably introduces noise and can prove inaccurate at times [6]. Furthermore, models trained on altered low-illumination images have not shown improved results over models only trained on low-illumination images [7]. While image preprocessing carries with it the potential to drastically enhance low-illumination object detection if done in real-time, this project has elected to focus on the capabilities of sensor fusion.

B. Improving energy conservation in autonomous vehicles

Research on optimizing energy efficiency in AVs has been diverse and effective, with many projects focusing on various components relevant to autonomous vehicles and detection algorithms. These works include applying the energy-efficient Spiking Neural Networks (SNN) to camera-based object detection, optimizing backbone networks for real-time detection algorithms, and assigning suboptimal, yet highly efficient detection algorithms to subsets of sensors in camera networks

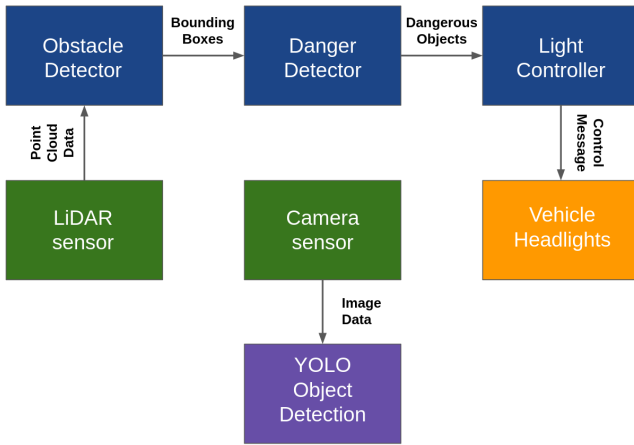


Fig. 1: ICanC system overview

[8]–[10]; however, to the best of the authors’ knowledge, no work has actively taken advantage of low-illumination to lower energy consumption in AVs. Autonomous vehicle research is prominent, and AVs will continue to improve in these areas.

III. ICAN C DESIGN

A. Overview

The ICanC system is designed to enhance vehicle safety by detecting obstacles, assessing potential dangers, and controlling vehicle lights to improve visibility. The system is composed of three interconnected nodes: the Obstacle Detector node, the Danger Detector node, and the Light Controller node. Each node performs specific functions that contribute to the overall functionality of the ICanC system.

The system overview in Fig. 1 illustrates the interactions and relationships between these nodes, providing a visual representation of the system’s operation. To utilize the system, only sensor data and control over headlights need to be given to ICanC; this promotes transferability and in theory, allows usage of ICanC across numerous platforms.

B. Methodology

The ICanC system comprises three nodes. Firstly, LiDAR sensors output point cloud data to the **Obstacle Detector** node, whose primary purpose is to identify objects and tightly fit 3D bounding boxes to each object. Next, these bounding boxes are sent to the **Danger Detector** node, which utilizes a variety of available data to determine whether or not each detected object has potential for danger to the vehicle. These dangerous objects are then collected and sent to the **Light Controller** node, which simply controls the light, allowing the camera sensors better imaging in low-illumination. These nodes are each discussed in full detail in the following subsections.

1) *Obstacle Detector Node*: There are numerous LiDAR object detection algorithms available for potential use in autonomous vehicles. The algorithm primarily used in the Obstacle Detector is Euclidean Clustering. It works as follows: the LiDAR sensors will formulate a point cloud frame, sending

this data to this node. The node will iterate over each point in this cloud, determining the Euclidean distance between other points in the cloud and clustering points within a set radius of the current point. Through this technique, a cluster of points returned from the sensors is grouped, allowing the node to tightly fit a box around the cluster’s edge points.

Additionally, this node calculates information relevant to the Danger Detector node. The Obstacle Detector keeps track of consecutive frames. With information concerning bounding boxes calculated, this node compares these consecutive frames and determines which objects are the same between frames. It does so by comparing overall displacement and intersection over union (IOU) between potential matches, collecting potential matches in a list, and then applying the Hungarian matching algorithm, assigning optimal matches to objects between frames.

With information from matching boxes identified, more important attributes relevant to danger detection can be calculated. Given two boxes representing the same object through consecutive frames, velocity can be determined in m/s as a calculation of translation between boxes divided by the time between the creation of the boxes. It should be noted that the node takes the translation of the vehicle as well as the translation of the object into consideration; however, the LiDAR sensor on the vehicle is the fixed frame, so the vehicle’s translation is simply applied to the object’s translation as well. Next, the orientation of the object can be calculated in the current frame through the angle between vectors formula:

$$\theta = \cos^{-1} \left(\frac{\vec{u} \cdot \vec{v}}{\|\vec{u}\| \cdot \|\vec{v}\|} \right),$$

where u and v are the x and y coordinates of the boxes $\in \mathbb{R}^2$ after translating the positions of the objects to the origin².

Early iterations of the ICanC system suffered from noisy and inaccurate sensor readings. Positions of bounding boxes were seemingly spontaneously assigned, leading to erroneous calculations. To combat faulty and inconsistent readings, a Kalman filter was implemented in this node to estimate the states of object positions and velocities in the 3-dimensional plane. In this filter, velocity is assumed to be constant. Previously calculated velocity and positions are fed to the filter, and readings are updated accordingly for orientation as well. Through filtering observed readings across frames, measurements are smoothed out, effectively combating unpredictable and noisy sensor data.

With all the necessary data calculated and filtered, an array of bounding boxes is sent to the Danger Detector node, each box carrying information on its position, velocity, and orientation.

2) *Danger Detector Node*: The initial design of the Danger Detector is shown in Fig. 2 and is as follows. The Danger Detector splits the frontal 180° field-of-view (FOV) of the vehicle into three distinct sections: the leftmost 45°, the center 90°, and the rightmost 45°. The logic for the danger detector

²Only yaw, or rotation about the z -axis, is considered in orientation calculation.

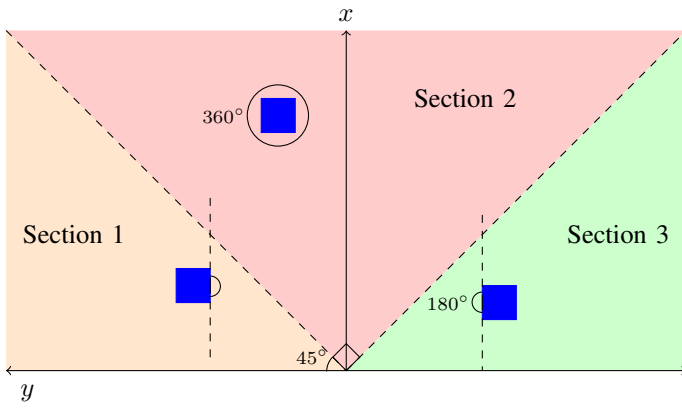


Fig. 2: Initial Danger Detector design

relies on mainly two factors: velocity and orientation. However, how these two are handled differently in the side and center sections.

Side sections: If an object is determined to be present in the side FOV of the vehicle, the orientation of the object is taken heavily into account. Specifically, if the object's previously calculated yaw faces the 180° toward the vehicle, the Danger Detector determines that the object is facing the vehicle. More specifically, if an object to the left of the vehicle is facing anywhere from 180° to 360° ³, the Danger Detector deems it facing the vehicle. For the rightmost FOV section, if it faces anywhere from 0° to 180° , it is facing the vehicle.

Next, the speed is taken into account. Given a float parameter t representing reaction time allowed, the distance at which the node deems a particular object dangerous is calculated. The system aims to allow for t seconds for the vehicle or passenger to react to the object after the light has been turned on. So, the distance formula $d = s \times t$ is used to calculate the approximate distance the current object will travel within t seconds.

Bringing these two calculations together, if the object's distance from the vehicle is currently less than the distance which the object will travel in t seconds and the object has been deemed facing the vehicle, the object has the potential to cause danger and the node labels the current object dangerous.

Front section: The danger of an object directly in the path of the vehicle is a more immediate issue. So in this case, the node disregards orientation altogether. Otherwise, the logic remains the same. If the distance of the object to the vehicle is less than the distance which the object will travel within t seconds, the Danger Detector deems the object dangerous.

There are a few flaws evident in this initial design, each contributing to the overall issue in the implementation: the danger label was applied indiscriminately to objects, outputting numerous false positives. Firstly, the center section covered a vast area, equal to half of the vehicle's FOV. Considering that the Danger Detector was considerably more lenient with objects found in this section, ignoring the orientation of the object

³Yaw orientation starts with 0° facing directly ahead of the vehicle, and 180° is facing directly behind the vehicle.

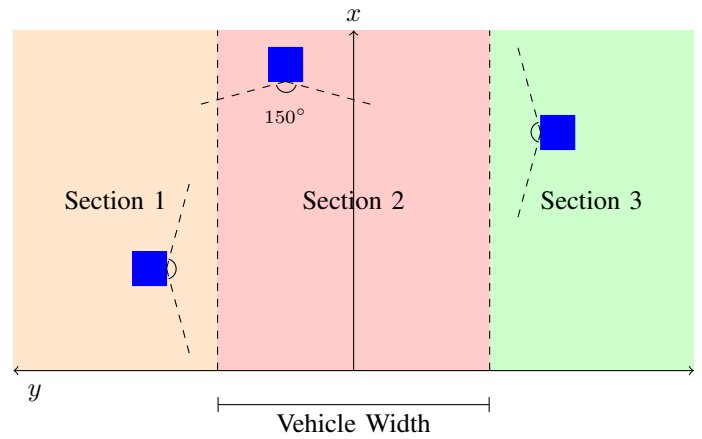


Fig. 3: Current Danger Detector design

altogether, multiple detected objects were given a dangerous label. To put this into perspective, Fig. 2 demonstrates the range of the center section. Specifically, an object directly ahead of the vehicle could be given the same level of danger as an object far to the left or right of the vehicle. In the side sections, where orientation is taken into account, the 180° FOV toward the vehicle allowed for objects was also overly lenient. In particular, an object could be moving parallel to the vehicle, never to collide, yet the Danger Detector would place the danger label on the object indefinitely. Overall, the initial design contained a few flaws, causing the system to designate far too many objects as dangerous.

The current design for the Danger Detector node is shown in Fig. 3 and improves upon these issues in a few ways. Firstly, the three sections are defined differently. Given a float parameter w representing the width of the vehicle, the Danger Detector splits the vehicle's frontal FOV into three sections: the direct path from $-w < y < w$, the leftmost FOV $y \geq w$, and the rightmost FOV $y \leq -w$. This adjustment decreases the overall area of the more dangerous direct path of the vehicle. Behavior for each section is defined below.

Side sections: The orientation allowed for objects detected in the side sections has been decreased from 180° to 150° . More specifically, for objects in the left FOV of the vehicle, if the Obstacle Detector node calculated the object's orientation to be between 195° to 345° , the Danger Detector labels it as facing the vehicle. For objects in the right section, the FOV allowed for danger is from 15° to 165° rather than 0° to 180° . This increases the criteria for dangerous objects and eliminates the issue of objects moving parallel to the vehicle being deemed dangerous. Otherwise, the logic remains the same.

Front section: Previously, the Danger Detector disregarded orientation altogether with the idea that an object within the direct path of the vehicle would be dangerous regardless of orientation. However, considering that the translation calculation from the Obstacle Detector takes into account the vehicle's movement and simply applies the vehicle's translation to the object, the orientation of the object should factor into the logic.

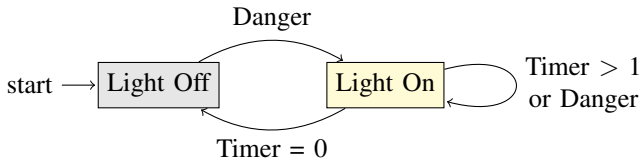


Fig. 4: State diagram for the Light Controller

For example, if an object had orientation 0° , or away from the vehicle, that means that the overall displacement of the object is away from the vehicle from one frame to the next. Since the object is moving away, it shouldn't be deemed dangerous. To combat this, logic from the side sections is implemented into the front section. The Danger Detector allows for 150° FOV toward the vehicle to be deemed facing the vehicle, from 105° to 255° . Other than the addition of orientation into the logic, the overall calculation remains the same for determining if an object is dangerous.

The Danger Detector collects all objects deemed dangerous into an array of bounding boxes and passes them into the Light Controller node.

3) *Light Controller Node*: The Light Controller has two states, shown in Fig. 4, in which it can reside: light on or off. Each of these states carries with it differing logic. When this node receives data from the Danger Detector in the off state, the Light Controller turns the light on. Given an integer parameter τ , the Light Controller will also start a timer for τ seconds, during which the light will continue to shine regardless of the presence or absence of danger. This time frame helps to eliminate any inconveniences that could arise from the light flashing on and off frequently. While the Light Controller is in the light-on state, it will ignore any new messages until the timer reaches one second. When the timer reaches one second, the light will still be on, but the Light Controller will begin to check whether or not the LiDAR still detects danger. If the Danger Detector passes new information to the Light Controller during this time frame, the timer will reset and the light will continue to shine for a new cycle. If no new information is received from the Danger Detector, the Light Controller will turn the light off and the timer will stay at zero. From here, the logic cycle repeats.

IV. IMPLEMENTATION

This system is built using the open-source Robot Operating System (ROS). Each node mentioned in the previous section is a ROS node, equipped with publishers and subscribers to communicate with one another through ROS topics. The LiDAR and camera object detection algorithms utilized were written with ROS wrappers, allowing for compatibility with the system. The ICanC system employs a Euclidean clustering LiDAR object detection algorithm as well as the single-pass YOLOv3 camera detection algorithm [11]–[13]. The Obstacle Detector also employs a C++ implementation of a Kalman filter [14], [15].

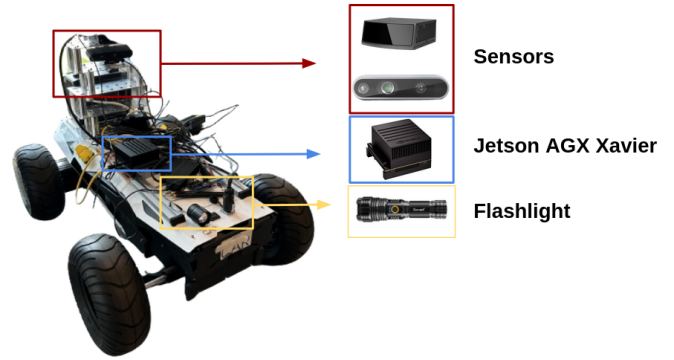


Fig. 5: Zebra: A general use indoor robot

Concerning the functions written for ICanC, the code and logic outlined in Section 3 can be abstracted and implemented into the Vehicle Programming Interface (VPI) to address a broader range of challenges in autonomous driving. Three key functions have been identified for their potential versatility: *getSection*, which determines whether an object is to the left, middle, or right of the vehicle; *isMovingTowardMe*, which assesses whether a detected object is moving toward the vehicle; and *getSpeedOfObject*, which calculates an object's speed. The pseudocode for ICanC's danger detecting logic implemented with the VPI is provided in Listing 1. This system is one such case where these functions prove useful, but these functions also may play critical roles in applications such as obstacle detection, collision avoidance, adaptive cruise control, and pedestrian tracking. By generalizing these functions into the VPI, both the efficiency and accessibility of autonomous vehicle programming is enhanced.

```

import vpi

def detect_danger(boxes):
    reaction_time_allowed = 3 #in seconds

    for box in boxes:
        speed = vpi.getSpeedOfObject(box)
        section = vpi.getSection(box)
        dist = norm(box) #dist from vehicle to
            object in meters

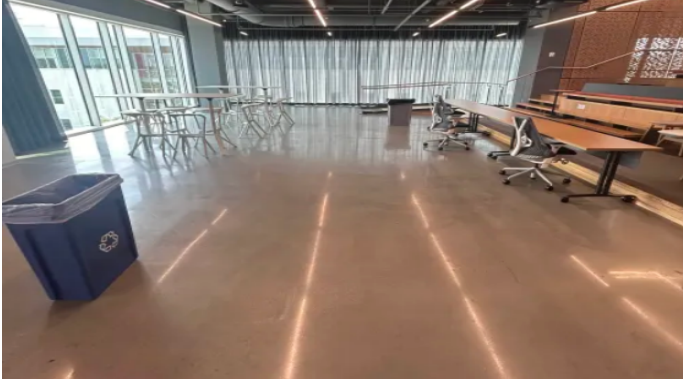
        danger_threshold = speed *
            reaction_time_allowed

        if (dist <= danger_threshold and
            vpi.isMovingTowardMe(box, section)):
            dangerous_boxes.append(box)

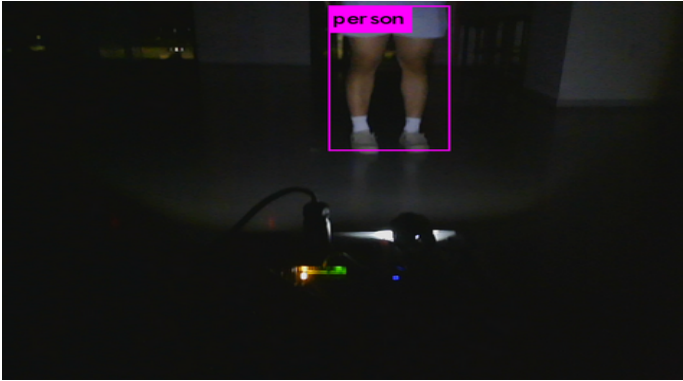
    return dangerous_boxes
  
```

Listing 1: Danger Detecting with VPI

This system was evaluated in three testing environments: a physical environment, a simulation environment via the CARLA autonomous driving simulator, and a Monte Carlo simulation.



(a) One physical environment used in testing.



(b) Camera object detection in physical environment

Fig. 6: Physical testing

A. Physical Testing

The ICanC system was installed onto the CAR Lab’s Zebra, a general-purpose indoor autonomous vehicle, and demonstrated in various lighting conditions and environments.

1) *Zebra*: The Zebra’s computing unit is the NVIDIA Jetson AGX Xavier Developer Kit running Ubuntu 18.04 and ROS Melodic [16]. This autonomous vehicle is shown in Fig. 5, equipped with a Robosense LiDAR M1 sensor as well as an Intel RealSense DepthCamera D435i, which features an RGB sensor with maximum resolution of 1920×1080 [17], [18]. Furthermore, the Zebra’s chassis is the AgileX hunter robot, which integrates an Ackermann control-based drive-by-wire [19]. This allows for both manual and autonomous control of the vehicle. To simulate headlights for the Zebra, a Goreit flashlight was interfaced with Arduino and attached to the vehicle [20], [21].

2) *Course/Setup*: Testing was conducted in the FinTech building at the University of Delaware at multiple sites. One such course is shown in Fig. 6a and 6b, arbitrarily populated with objects, such as garbage bins, chairs, and tables, to challenge the ICanC system. The testing took place at nighttime, with building lights on and off. The tests were also run with and without the flashlight attached to the Zebra to highlight the differing capabilities of LiDAR and camera sensors in the dark. Furthermore, a pedestrian was introduced to the scene,

moving toward, away from, and perpendicular to the Zebra in different tests.

B. Simulation Testing

The CARLA simulator provides realistic environments with control over the time of day, sensors, streetlights, and headlights, allowing accurate simulation of environments where the ICanC system would be effective [22].

1) *Vehicle*: The vehicle equipped with ICanC in the simulation is the Tesla Model 3. This vehicle is armed with a custom LiDAR sensor with 128 channels and 640,000 points per second generated, as well as an RGB camera sensor. Parameters were adjusted to accommodate the Tesla’s width w and overall FOV of the LiDAR to allow for a wider radius of detection than the smaller Zebra. Fig. 7a and 7b shows the vehicle in the pitch-black environment, as well as its LiDAR and camera detection outputs in Fig. 7c and 7d. This vehicle was equipped with both autonomous and manual control capabilities, and both were employed to observe the behavior of the system.

2) *Environment*: Testing was done in various preloaded environments offered by CARLA, including Town01. To achieve the desired world/environmental state, scripts were run to change the time of day to midnight, set all streetlights to off, and spawn numerous cars, trucks, and pedestrians in random areas. This created a pitch-black environment, as shown in Fig. 7.

From there, ICanC was activated, allowing point cloud and camera data to be fed into the system’s nodes.

C. Monte Carlo Simulation

The goals of this Monte Carlo simulation are to:

- Determine the relationship between the number of objects detected in a frame and the percent of those objects deemed dangerous.
- Determine the average percentage of objects that the Danger Detector deems dangerous.
- Utilize this information to calculate approximate energy conservation.

The hypothesis before running the simulation was that the distribution would be uniform and that there was no special relationship between load size, or number of objects detected, and percent danger detected.

To support this hypothesis, the probability of danger was calculated beforehand. Randomly generated bounding boxes representing detected objects were created with the following random variables: distance d pulled from a discrete uniform distribution $\text{unif}\{1, n\}$ in meters, speed s pulled from the discrete uniform distribution $\text{unif}\{1, n\}$ in meters per second, and orientation ψ pulled from a discrete uniform distribution $\text{unif}\{1, 360\}$ in degrees. The probability of a detected object being deemed dangerous is as follows.

Let θ be the danger distance threshold, defined as

$$\theta = \alpha \cdot s,$$

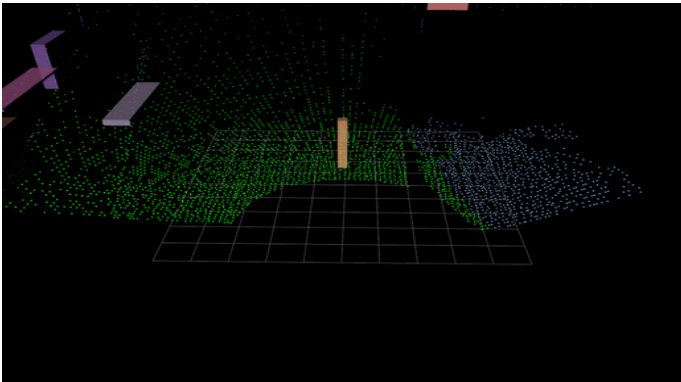
where α represents the parameter for reaction time allowed for the vehicle. In this simulation, $\alpha = 1$, leaving us with $\theta = s$.



(a) ICanC equipped vehicle approaching another vehicle in the dark.



(b) Other vehicle deemed as dangerous, headlights activate



(c) Obstacle Detector detects an object directly ahead of sensors.



(d) Object is deemed dangerous, camera object detection identifies it as a traffic light.

Let E_1 be the event that $d \leq \theta$, or that the current position of the detected object is less than the danger threshold calculated in θ , and let E_2 be the event that the detected object is facing the vehicle. The event space Ω is defined through the two events and their four total combinations. Finally, let D be the event that the detected object is dangerous.

Assuming that E_1 and E_2 are independent variables, the probability of D can be defined as

$$P(D) = P(E_1) \wedge P(E_2)$$

In the above equation, $P(E_2)$ can be defined as $\frac{150}{360} \approx .4166$, as this is the allowed FOV to be deemed facing the vehicle.

For $P(E_1)$, defined through the cumulative distribution function $P(E_1) = P(x \leq \theta)$, the probability of each discrete point from 1 to n needs to be considered. Let event C_i be the event that the object is i meters away from the vehicle, where $i \in [1, n]$. $P(E_1)$ can be expanded using the law of total probability.

$$\begin{aligned} P(E_1) &= \sum_{i=1}^n P(E_1|C_i) \cdot P(C_i) \\ &= \sum_{i=1}^n \frac{i}{n} \cdot \frac{1}{n} \\ &= \sum_{i=1}^n \frac{i}{n^2} \\ &= \frac{1}{n^2} \cdot \frac{n(n+1)}{2} \\ &= \frac{n^2 + n}{2n^2} \\ &= \frac{1}{2} + \frac{1}{2n} \end{aligned}$$

We can further conclude that as n , or the range of the LiDAR sensors increases,

$$\lim_{n \rightarrow \infty} \frac{1}{2} + \frac{1}{2n} = \frac{1}{2},$$

though LiDAR sensors usually will not detect objects further than a few hundred meters away.

Taking these values and introducing them into the equation for the probability of danger, the outcome is

$$P(D) = \frac{5}{12} \cdot \frac{1}{2} = \frac{5}{24} \approx .2083.$$

The simulation operates under the same assumptions as the previous calculations, with n set to 60. With $n = 60$, $P(E_1) \approx .5083$, and $P(D)$ is calculated to be .2118. To determine the accuracy of this calculation and the previous hypothesis, Algorithm 1 was formulated.

Running this algorithm outputs the results in Fig 8, which will be discussed further in the evaluation.

It should be noted that this simulation operates under assumptions that are not entirely realistic, as objects in physical environments are rarely random. For instance, the vast majority of vehicles will follow a predictable path, adhering to traffic

Fig. 7: CARLA simulator testing

Algorithm 1 Monte Carlo simulation for ICanC system

Input: n (Highest number of boxes generated)**Output:** *Averages*

```
1: Averages  $\leftarrow$  []
2: for  $i = 1$  to  $n$  do
3:    $sum \leftarrow 0$ 
4:    $num\_trials \leftarrow 1000$   $\triangleright$  Monte Carlo trials
5:   for  $j = 1$  to  $num\_trials$  do
6:      $boxes \leftarrow \text{GenerateBoxes}(i)$   $\triangleright i$  boxes
7:      $\text{PublishBoxes}(boxes)$   $\triangleright$  Pass boxes into system
8:      $pct\_danger \leftarrow \text{GetPercentDanger}(i)$   $\triangleright$  Determine
       danger percentage given  $i$  boxes published
9:      $sum \leftarrow sum + pct\_danger$ 
10:  end for
11:   $average \leftarrow sum/num\_trials$ 
12:  Append  $average$  to Averages
13: end for
14: return Averages
```

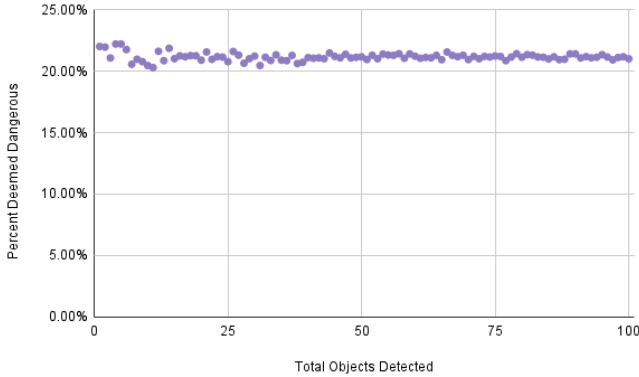


Fig. 8: Results from Monte Carlo Simulation on probability that the Danger Detector will label an object dangerous.

laws and lane lines. Nonetheless, this simulation offers insight into the predicted behavior of the Danger Detector as well as an estimate of the amount of danger detected by the node.

V. EVALUATION

The ICanC system was extensively tested in numerous environments, and was further evaluated on multiple metrics: GPU/CPU usage, latency, and energy conservation, along with the standard performance metrics.

A. Performance Metrics

Data was recorded during physical testing in the form of rosbag files, a format for recording ROS messages such as point cloud and image data. Each bag encompassed a unique scenario, where a pedestrian would either start close to or far from the vehicle. The pedestrian would then move toward, away from, or perpendicular to the vehicle. Each scenario was predetermined to have objects deemed as truly dangerous or

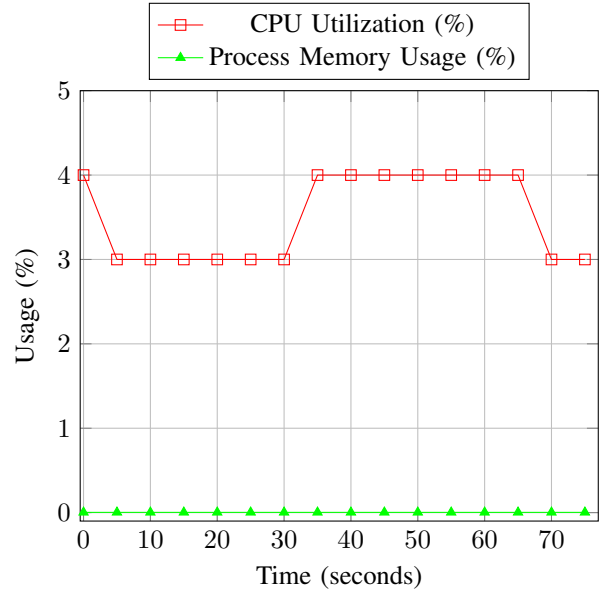


Fig. 9: CPU usage plotted against time

TABLE II: Performance Metrics

Metric	Value
Precision	57.14%
Recall	72.73%
Accuracy	74.29%
F1 Score	64.00%

not dangerous. True positives, true negatives, false positives, and false negatives were then determined, and the performance metrics in Table II were calculated.

The metric most relevant to ICanC is recall, which weighs true positives and false negatives. ICanC controls a light that assists camera object detection; there are negligible consequences for turning on the light more often than needed. However, failing to detect danger and outputting false negatives pose a substantial threat to the vehicle and its passengers. The recall score of the model is an acceptable score of 72.73%.

B. System Utilization

GPU utilization, CPU utilization, and memory usage were evaluated through the rosbags recorded in the physical environment. The system was run on a computer equipped with a Nvidia GeForce RTX 2080 GPU, and the system utilization was recorded every five seconds. The entirety of ICanC, including sensors and object detection algorithms, proved stressful on processing units; however, it can be reasoned that LiDAR and camera object detection algorithms, if equipped on an autonomous vehicle, will be initiated regardless. Thus, Fig. 9 displays the system utilization when ICanC nodes are isolated from detection algorithms. Overall, it can be concluded that the ICanC system doesn't subject the CPU to intense loads, only taking approximately 4% of the CPU during use and 0.003% of process memory, or 0.96 MiB. Furthermore, the system

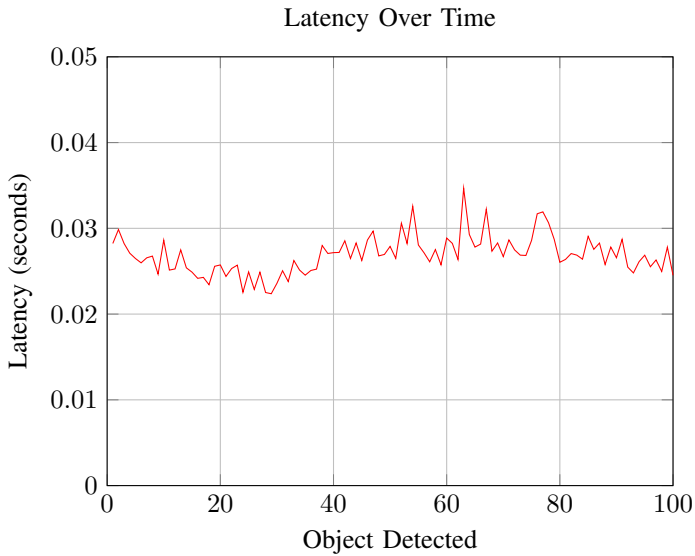


Fig. 10: Latency over time for 100 objects detected

doesn't utilize GPU at all, showing that the ICanC system is exceptionally resource-efficient.

C. Latency

End-to-end latency is recorded through rosbag data from the moment point cloud data is received by the Obstacle Detector to the moment the flashlight receives the command to turn on. The first 100 messages sent were measured and displayed in Fig. 10.

From receiving data to turning on the lights, the latency hovers around 30 ms. This time is well within the bounds of efficient computation given the context of AVs and real-time systems. A latency of 30 ms generally allows AVs or drivers ample time to adjust maneuvering given a dangerous object has been detected. Along with the parameter for reaction time allowed (during testing, three seconds were allowed to react after the light turned on), there should be sufficient time for response.

D. Energy Conservation

Energy usage can be measured in terms of wattage · hours of usage. Table III shows the watts used by various relevant lights, including the standard halogen headlights, newer LED headlights, and the Goreit tactical flashlight utilized in the physical testing environment [23], [24].

To evaluate the potential energy saved for each of these lights, a baseline scenario is considered: a 22-minute drive from the CAR Lab building at the University of Delaware in Newark, DE to Wilmington, DE near midnight. For this scenario, the worst case, average case, and best case for the ICanC system are considered. The system was not run during this drive; instead, energy conserved is calculated for each scenario assuming the system operates precisely how it was designed to operate. The calculations assume usage of high beams, as relatively standard during nighttime (see Fig 11).

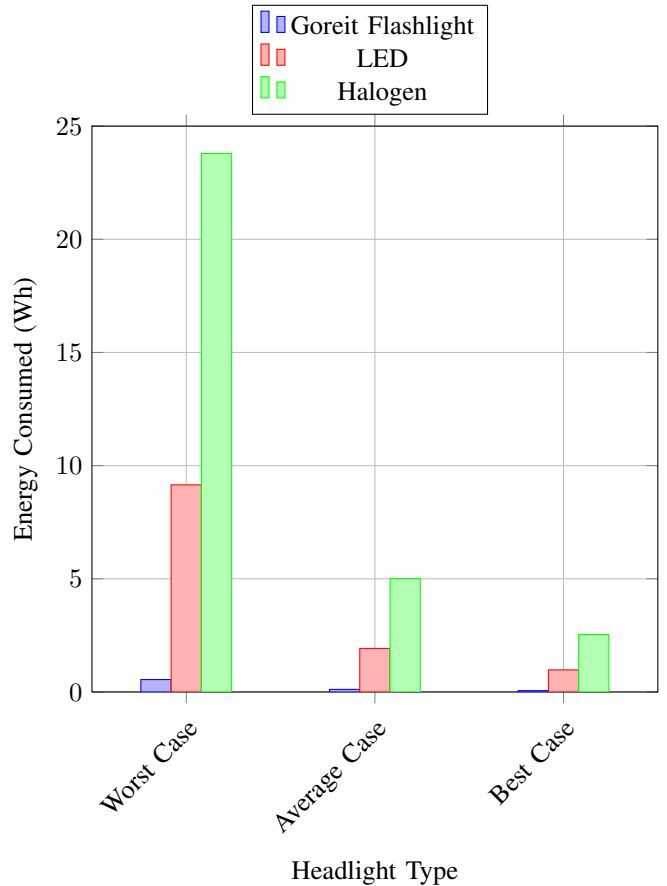


Fig. 11: Energy consumed in the worst case, average case, and best case scenarios for different headlight types.

Headlight Type	Low Beam (Watts)	High Beam (Watts)
LED	15	25
Halogen	55	65
Goreit Flashlight	N/A	1.5

TABLE III: Approximate wattage of different headlight types on low and high beam settings

1) *Worst Case*: The worst case is defined simply as having headlights on for the entire duration of the drive. This case can be considered the baseline as full usage of headlights is standard in low-illumination environments. By defining the worst case as the baseline, it's ensured that the ICanC system will provide energy-conserving results.

A quick calculation shows that $1.5 \text{ watts} \cdot 0.366 \text{ hours} = 0.549 \text{ Wh}$ consumed for the flashlight utilized in the system. For halogen and LED headlights, 23.79 Wh and 9.15 Wh are consumed, respectively.

2) *Average Case*: Calculations for the probability that a detected object will be deemed dangerous as well as a dedicated Monte Carlo simulation have shown that approximately 21% of objects will be deemed dangerous (see Fig. 8). Note that, once again, the simulation and the calculations operate under

inaccurate assumptions, but assuming that the system detects a random object every second, this data offers an average case benchmark.

The light will be on for 4 minutes and 37 seconds, or .077 hours. Given this, the flashlight, halogen, and LED lights expend 0.1155 Wh, 5.005 Wh, and 1.925 Wh respectively.

3) *Best Case*: The ideal case is based on personal observations on the drive from Newark to Wilmington. During the drive, objects believed to be potentially dangerous were counted. These objects included traffic signs and cones close to the vehicle, pedestrians walking on the sidewalk, and other vehicles driving nearby. In the 22-minute drive, 47 objects were personally deemed dangerous.

Assuming the light timer parameter τ is set to three seconds, and each object was detected greater than 3 seconds apart, the light should be on for 141 seconds, or .039 hours. The flashlight, halogen, and LED lights would expend 0.0585 Wh, 2.535 Wh, and 0.975 Wh respectively.

Taking a look at each of these scenarios, an average case defined by the Monte Carlo simulation only uses 21% of the energy that the standard case uses, saving 7.225 Wh for an LED headlight, which AVs are often outfitted with. Furthermore, the best case uses approximately 10% of the standard case, saving 8.175 Wh for an LED headlight. While the energy savings from minimizing headlight usage may be relatively modest compared to the total energy consumption of the vehicle, it remains a significant consideration. Over time, these incremental savings can accumulate, and any reduction in energy usage is advantageous.

VI. DISCUSSION

In this paper, the ICanC system is formulated and implemented. The results suggest efficient processing utilization and speed, acceptable danger-detecting performance metrics, as well as the potential for significant energy conservation improvements in vehicle headlight usage. These findings imply applicability in autonomous vehicles, reducing energy usage without sacrificing accuracy or processing power.

It's important to note that the focus of this project is on low-illumination states. This project has not been tested in other adverse weather environments that would negatively affect camera-based object detection, including fog, rain, and snow. Most likely, ICanC would perform poorly in these conditions due to LiDAR backscattering. However, as sensor technology continues to evolve, it may become possible to achieve ICanC-esque functionality using weather-independent sensors.

In the future, improvements to ICanC can be achieved by modifying the LiDAR object detection algorithm. Several performance metrics, such as precision and F1 score, were subpar for danger detection. Observations revealed that the object detection algorithm based on Euclidean clustering was quite inaccurate and susceptible to noise. Although the Kalman filter significantly enhanced the accuracy and precision of the tracking algorithm, an improved object detection algorithm would further boost these metrics, one potential avenue being a trained Point Pillars model [25].

VII. CONCLUSION

In conclusion, ICanC was designed and implemented to promote energy conservation while maintaining accurate camera-based object detection. Extensive testing in real-world and simulated environments demonstrated its computational efficiency, speed, and potential to significantly reduce headlight energy consumption without compromising detection performance. Future work includes generalizing the system to address broader autonomous vehicle challenges and refining the LiDAR detection algorithm. Overall, ICanC successfully achieves its objectives, contributing to the advancement of sustainable transportation.

VIII. ACKNOWLEDGEMENTS

I would like to express my gratitude to the National Science Foundation for sponsoring this summer research experience at the University of Delaware and for placing me in a position to succeed.

REFERENCES

- [1] T. Thadani, R. Lerman, I. Piper, F. Siddiqui, and I. Uraizee, "The Final 11 Seconds of a Fatal Tesla Autopilot Crash," Washington Post, 10 2023. [Online]. Available: <https://www.washingtonpost.com/technology/interactive/2023/tesla-autopilot-crash-analysis/>
- [2] G. Zamanakos, L. Tsochatzidis, A. Amanatiadis, and I. Pratikakis, "A comprehensive survey of LIDAR-based 3D object detection methods with Deep learning for autonomous driving," *Computers & Graphics*, vol. 99, 07 2021.
- [3] B.-F. Wu, R. Zhong, Y. Wang, J. Wan, J.-L. Zhang, and W. Shi, "VPI: Vehicle Programming Interface for Vehicle Computing," *Journal of Computer Science and Technology*, vol. 39, pp. 22–44, 01 2024.
- [4] H. Patel and K. P. Upla, "Night Vision Surveillance: Object Detection using Thermal and Visible Images," *IEEE International Conference for Emerging Technology (INCET)*, p. 1–6, 06 2020. [Online]. Available: <https://ieeexplore.ieee.org/document/9154066>
- [5] F. Julca-Aguilar, J. Taylor, M. Bijelic, F. Mannan, E. Tseng, and F. Heide, "Gated3D: Monocular 3D Object Detection From Temporal Illumination Cues," 02 2021.
- [6] Y. P. Loh, X. Liang, and C. S. Chan, "Low-light image enhancement using Gaussian Process for features retrieval," *Signal Processing: Image Communication*, vol. 74, pp. 175–190, 05 2019.
- [7] Y. Xiao, A. Jiang, J. Ye, and M.-W. Wang, "Making of Night Vision: Object Detection Under Low-Illumination," *IEEE Access*, vol. 8, pp. 123 075–123 086, 2020.
- [8] S. Kim, S. Park, B. Na, and S. Yoon, "Spiking-YOLO: Spiking Neural Network for Energy-Efficient Object Detection," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, pp. 11 270–11 277, 04 2020.
- [9] Y. Lee, J.-w. Hwang, S. Lee, Y. Bae, and J. Park, "An Energy and GPU-Computation Efficient Backbone Network for Real-Time Object Detection," *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, p. 752–760, 06 2019. [Online]. Available: <https://ieeexplore.ieee.org/document/9025413/>
- [10] T. Dao, K. Khalil, A. K. Roy-Chowdhury, S. V. Krishnamurthy, and L. Kaplan, "Energy Efficient Object Detection in Camera Sensor Networks," *IEEE International Conference on Distributed Computing Systems (ICDCS)*, 06 2017.
- [11] S. Sun, "LiDAR Obstacle Detector," GitHub, 12 2021. [Online]. Available: https://github.com/SS47816/lidar_obstacle_detector
- [12] J. Redmon and A. Farhadi, "YOLOv3: An Incremental Improvement," *IEEE Computer Vision and Pattern Recognition (CVPR)*, 04 2018.
- [13] M. Bjelonic, "YOLO ROS: Real-Time Object Detection for ROS," GitHub, 12 2016. [Online]. Available: https://github.com/leggedrobotics/darknet_ros
- [14] Y. Pei, S. Biswas, D. S. Fussell, and K. Pingali, "An elementary introduction to Kalman filtering," *Communications of the ACM*, vol. 62, pp. 122–133, 10 2019. [Online]. Available: <https://arxiv.org/pdf/1710.04055.pdf>

- [15] H. Martiros, "Kalman C++," GitHub, 11 2014. [Online]. Available: <https://github.com/hmartiro/kalman-cpp>
- [16] "Nvidia Jetson Xavier Series," NVIDIA. [Online]. Available: <https://www.nvidia.com/en-us/autonomous-machines/embedded-systems/jetson-xavier-series/>
- [17] "RoboSense LiDAR M1," Robosense.ai. [Online]. Available: <https://www.robosense.ai/en/RS-LiDAR-M1>
- [18] "Intel® RealSense™ Depth Camera D435i," Intel® RealSense™ Depth and Tracking Cameras. [Online]. Available: <https://www.intelrealsense.com/depth-camera-d435i/>
- [19] "Hunter Chassis," Agilex Robotics, 2024. [Online]. Available: <https://global.agilex.ai/products/hunter-se>
- [20] "Goreit Flashlight," Amazon.com. [Online]. Available: <https://www.amazon.com/Flashlights-Rechargeable-Goreit-Flashlight-Emergency/dp/B09QMJ4D4J?th=1>
- [21] "Arduino Uno," Arduino. [Online]. Available: <https://store.arduino.cc/products/arduino-uno-rev3?srltid=AfmBOopl4jNQ6ZBL3-LkXBfgFEuvTBjegXKXhizfKkmTH0LacN8B61AOc>
- [22] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "CARLA: An Open Urban Driving Simulator," *Proceedings of Machine Learning Research (PMLR)*, vol. 78, 11 2017.
- [23] "Philips Halogen VisionPlus," Philips. [Online]. Available: <https://www.usa.philips.com/c-p/9007VPB2/visionplus-upgrade-headlight-bulb>
- [24] "Ultinon Pro6000 LED," Philips. [Online]. Available: <https://www.philips.com/c-e/au/road-legal-led-headlights/ultinon-pro6000-led.html>
- [25] A. H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom, "PointPillars: Fast Encoders for Object Detection From Point Clouds," *IEEE Computer Vision and Pattern Recognition (CVPR)*, p. 12689–12697, 06 2019. [Online]. Available: <https://ieeexplore.ieee.org/document/8954311>