

Camera-based Automotive License Locator: CALL

F6: Eric Tang, Richard Sbaschnig, Tzen-Chuen Ng
18-500 Capstone Design, Spring 2025
Electrical and Computer Engineering Department
Carnegie Mellon University

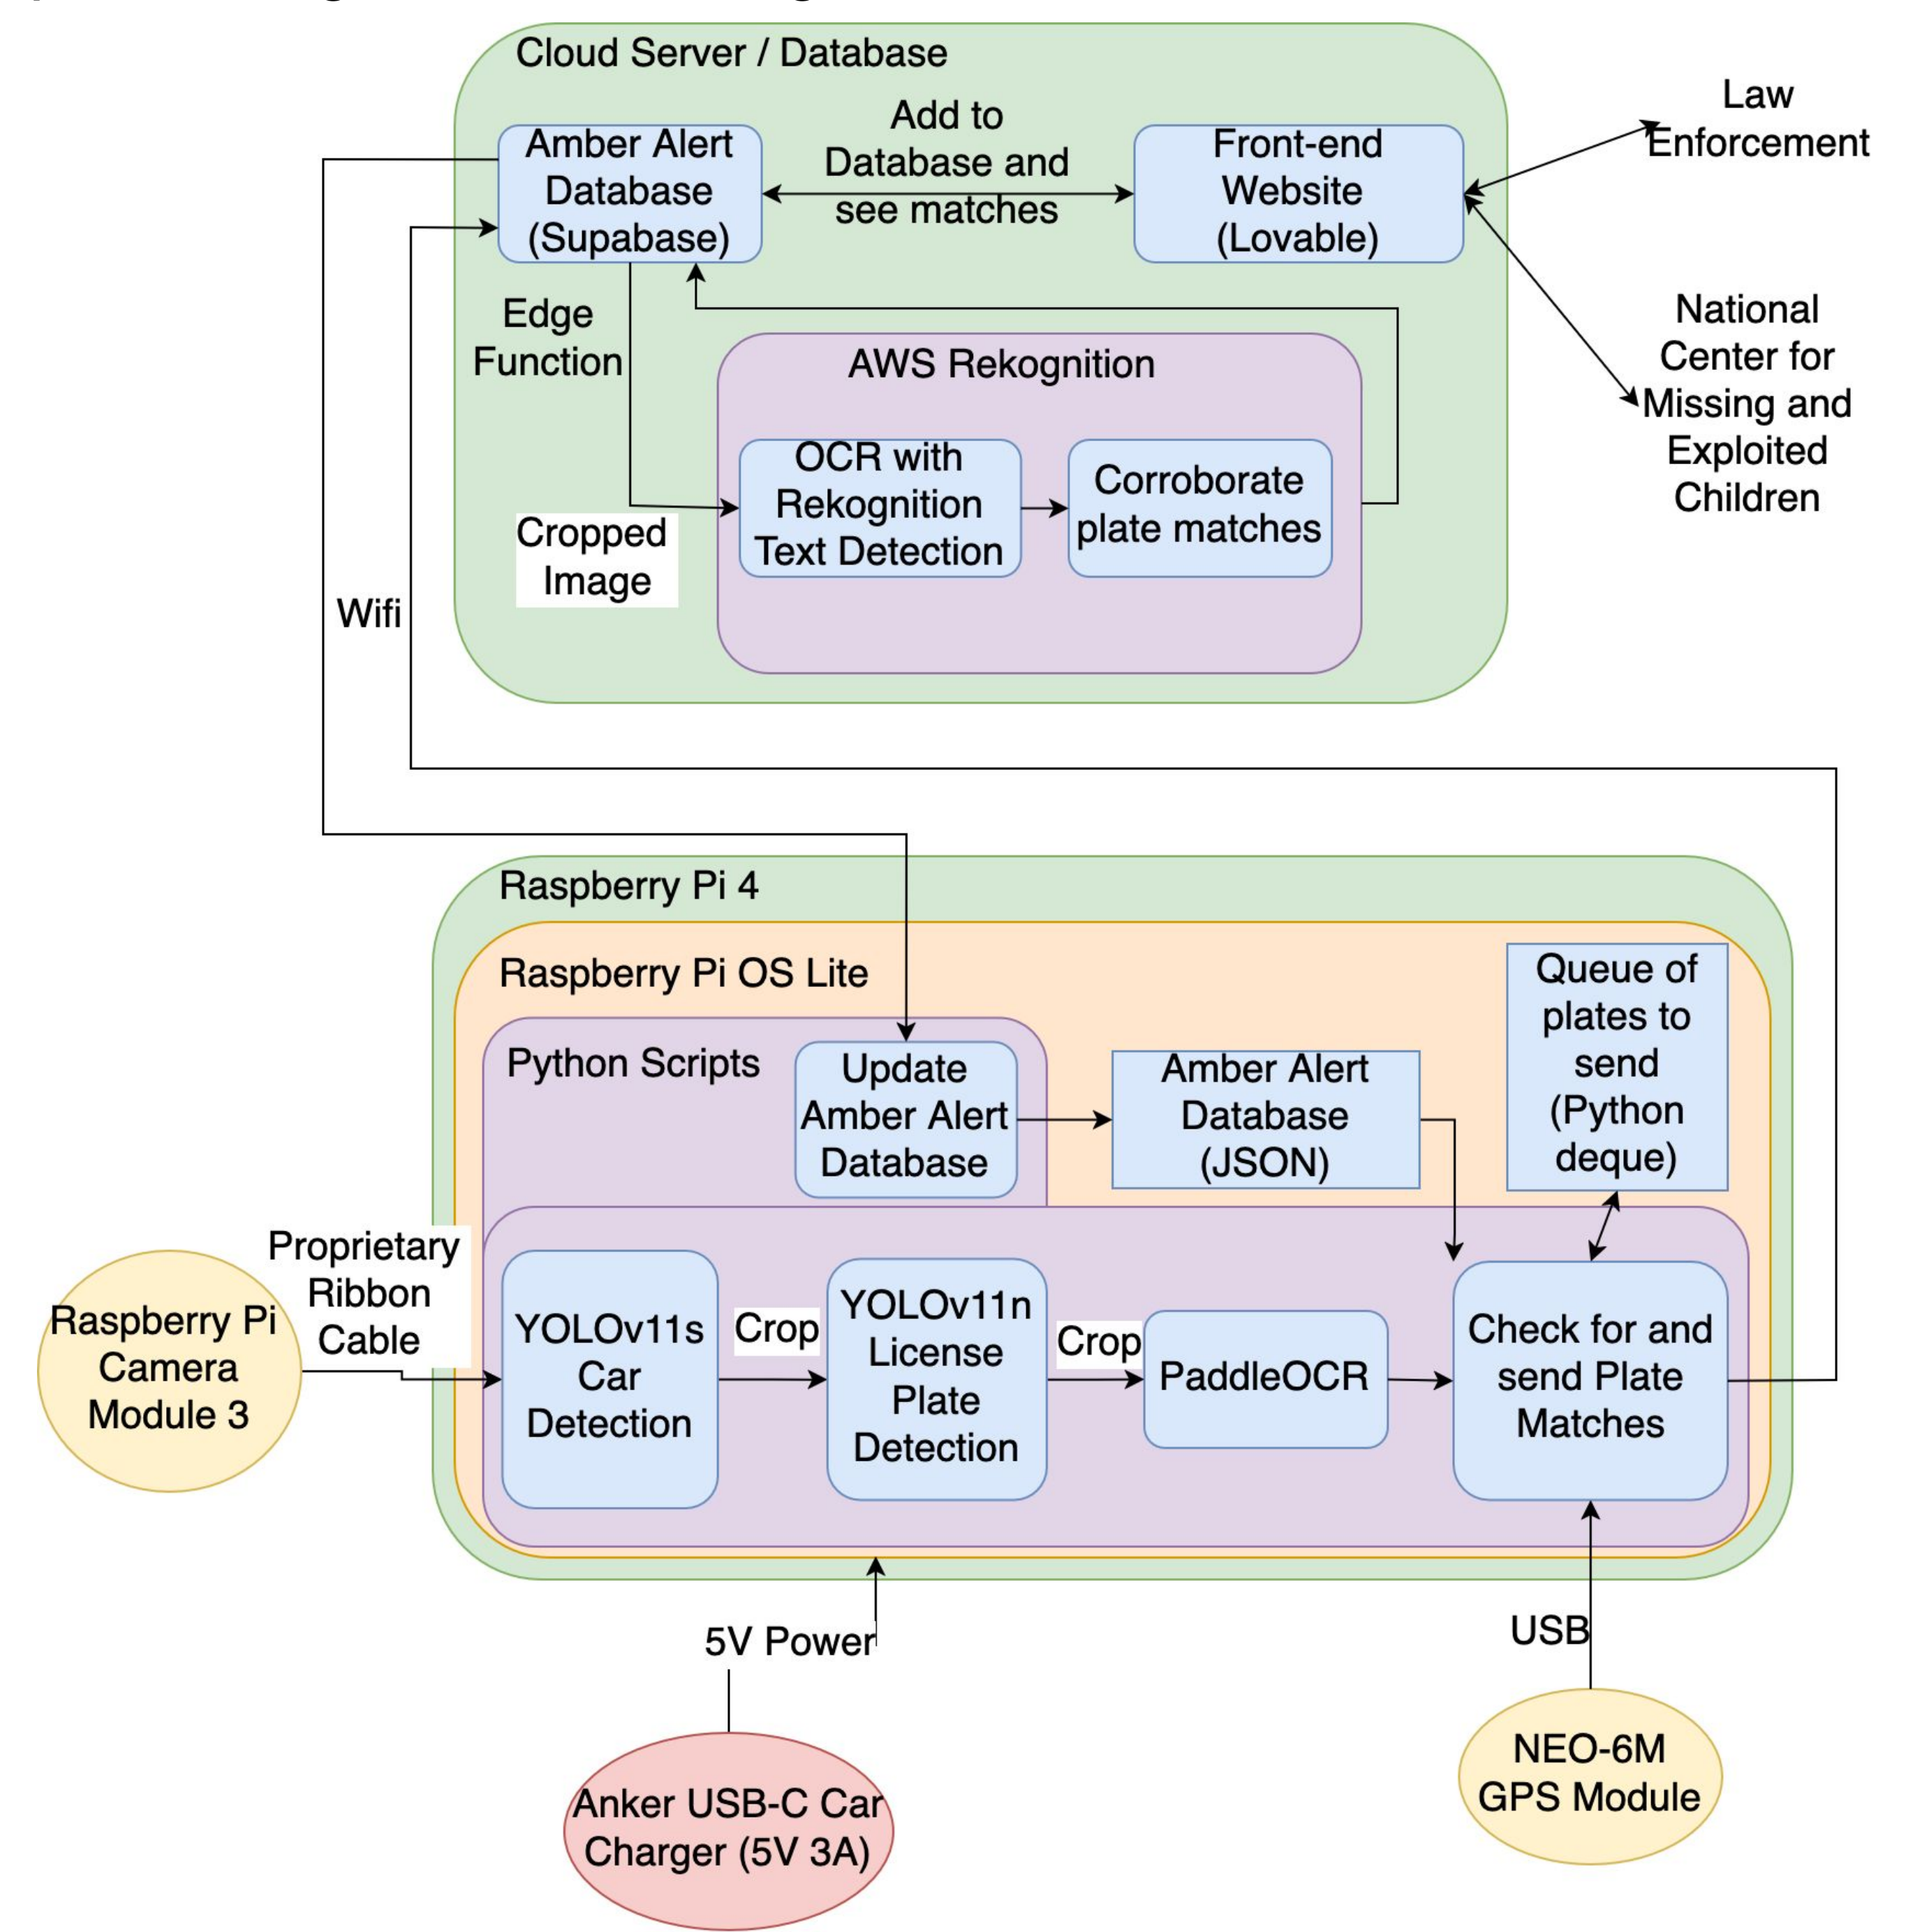
Product Pitch

In 2023, over 375,000 children were reported missing in the U.S., but only a fraction triggered Amber Alerts — and even fewer were solved because of them. To expand the effectiveness of Amber Alerts without overwhelming the public, we created CALL, a low-cost, dashcam-style device that passively scans license plates to identify suspect vehicles.


Our system scans the roads ahead **every 40 seconds** to seamlessly recognize license plates in at least a **10 meter range** ahead of our vehicle with a **91.9% precision and 84.7% recall**. The matched license plates are then sent to a database in the cloud - updated by Amber Alert officials - along with GPS location accurate **within 31m** for law enforcement to review.

System Architecture

A Raspberry Pi unit running our machine learning models captures and processes license plates locally, then sends potential matches to a cloud server where further cloud processing verifies and logs confirmed matches.



Conclusions & Additional Information




CALL is a compact, low-cost system that detects suspicious license plates and passively supports Amber Alert efforts, demonstrating a scalable approach to enhancing child recovery efforts. Throughout the project, we learned how to balance edge and cloud resources to optimize both performance and privacy. We also gained experience integrating real-time machine learning pipelines on resource-constrained hardware, handling unreliable input data from dynamic environments, and designing a system that prioritizes ease of deployment for everyday users. Future improvements could include expanding vehicle attribute detection and further fine-tuning the edge ML models for better real-world performance.

System Description

CALL can be organized into three distinct subsections: The camera, the edge compute unit, and the cloud server. Our chosen processor is the Raspberry Pi 4, which paired very nicely with the Raspberry Pi Camera Module 3 for our camera. Running on the Raspberry Pi is the PaddleOCR optical character recognition model, and two YOLOv11 object detection models, one for detecting cars and one for detecting license plates. The cloud server is Supabase, a PostgreSQL database that centrally holds relevant data such as the Amber Alert database and received matches, and is connected to AWS Rekognition for server-side match verification.

The workflow is as follows: An official enters amber alert information into the database, which is then loaded onto each CALL system. The camera captures images at intervals, which are processed locally, and if matched, is sent to the cloud for further verification. Once verified, the image and GPS information are posted for law enforcement to act upon.



Website and Hardware Setup (left)
Machine learning models (right)

System Evaluation

Use-Case/Design requirements:

Metric	Target	Actual
Timing Target	Image processed every 40 seconds	Met target
ML Model Precision and Recall	90% Precision 85% Recall	98.6% Precision 79.1% Recall
GPS Accuracy	Within 200m	Within 31m
Camera Quality	90% Precision 85% Recall	100% Precision 88.9% Recall
Edge-Cloud Communication	Updates properly 100% of the time	Met target
System Precision and Recall	90% Precision 85% Recall	91.9% Precision 84.7% Recall

System Trade-Off Factors:

- Edge vs Cloud Inference
- Camera Quality vs Cost
- Model Size vs Device Constraints