Carnegie Mellon University

F6: CALL

Eric Tang, Tzen-Chuen Ng, Richard Sbaschnig

Use Case and Design Requirements

Simple
Cost effective and easy to install

→ Runs off car socket
→ Under $100

Accurate
Reliably captures suspected plates

→90% precision
→85% recall

Fast
Captures all plates in front of car frequently

→40s snapshot
→Scalable cloud server

Robust
Operates at 10m range in adverse weather and lighting

→Autofocus
→Infrared
→High Image Quality
→High FOV

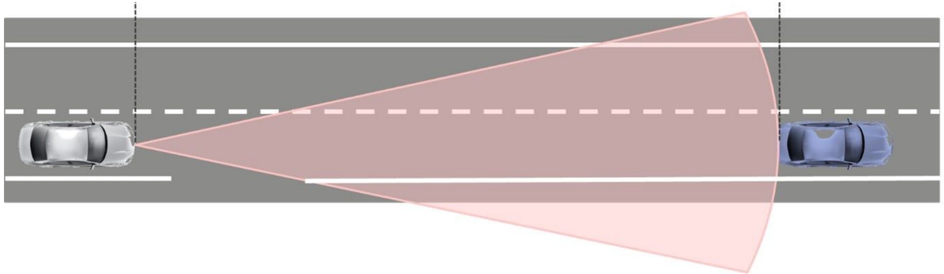

Solution Approach

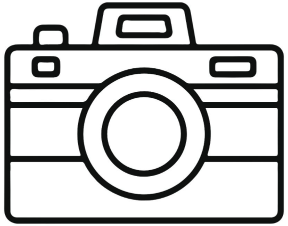

Camera Module 3
High resolution images with IR capabilities

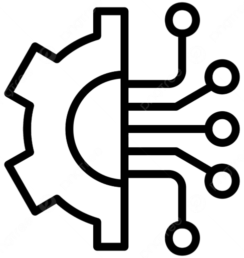

Raspberry Pi 4, Paddle OCR, YOLOv11
Low cost small-footprint hardware running a lightweight ML system

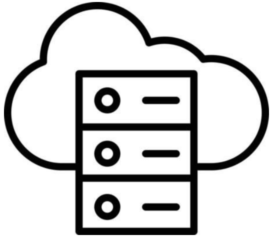

Supabase, Rekognition
PostgreSQL database connected to more powerful ML models for verification, storage, and access.

Security and Privacy considerations
Segmented, multi-user platform to contain sensitive location data

Solution Approach: Block Diagram

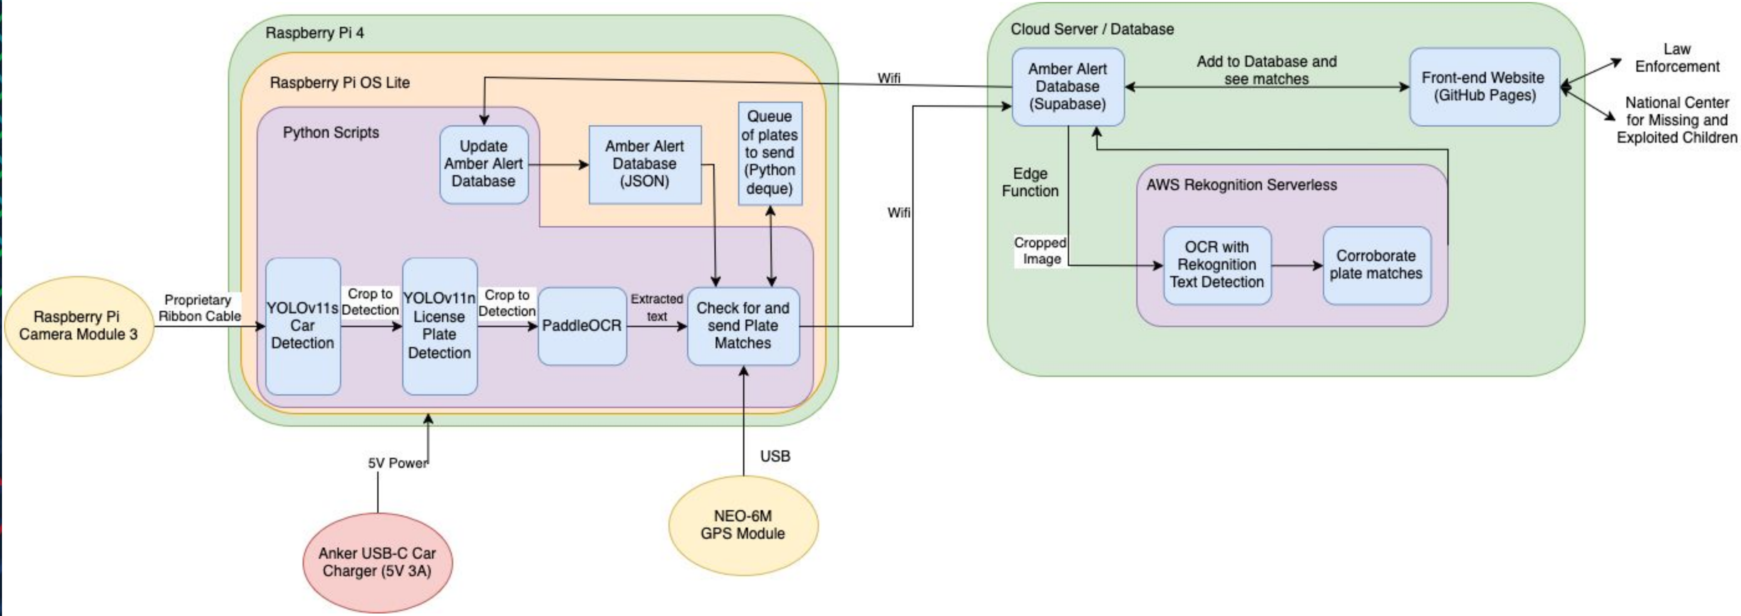

Complete Solution: User Approach

Amber Alert Personnel

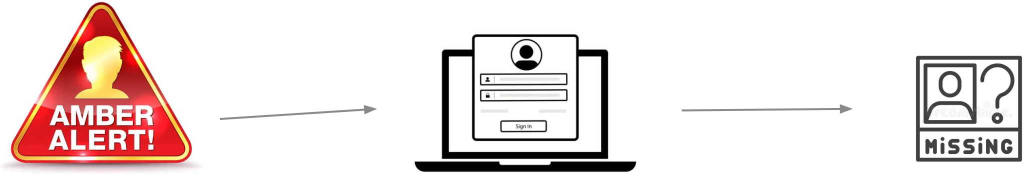

Law enforcement

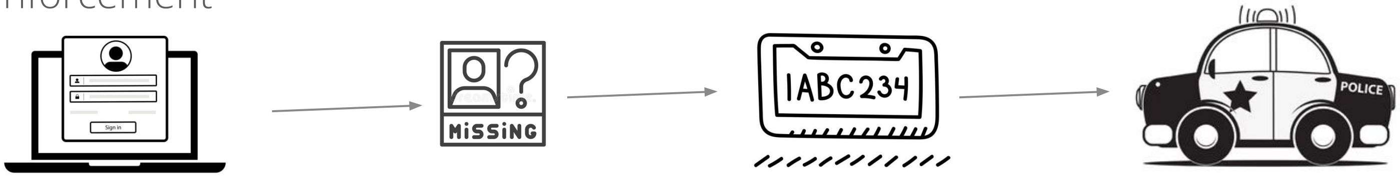

Dashcam Users

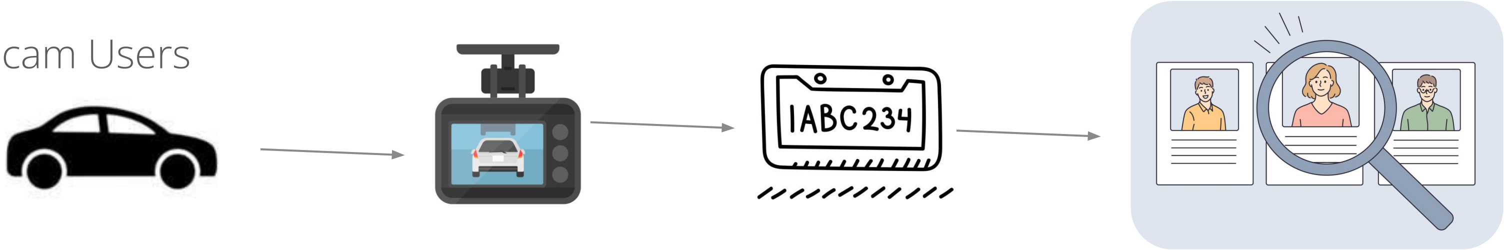

Simulated with video
from out camera

Carnegie Mellon University

Complete Solution: User Approach

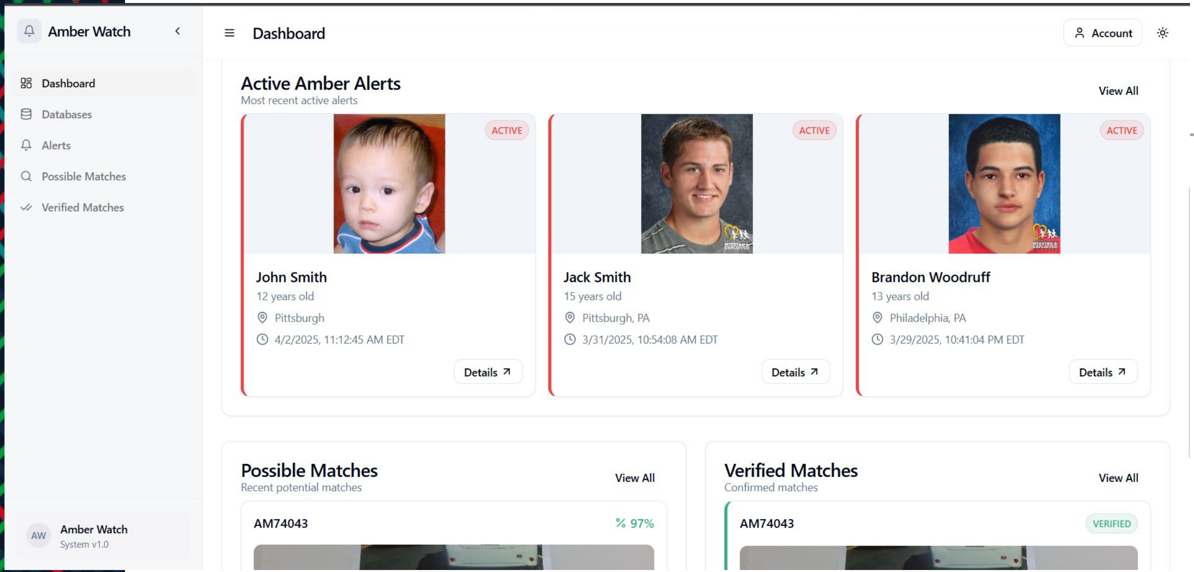

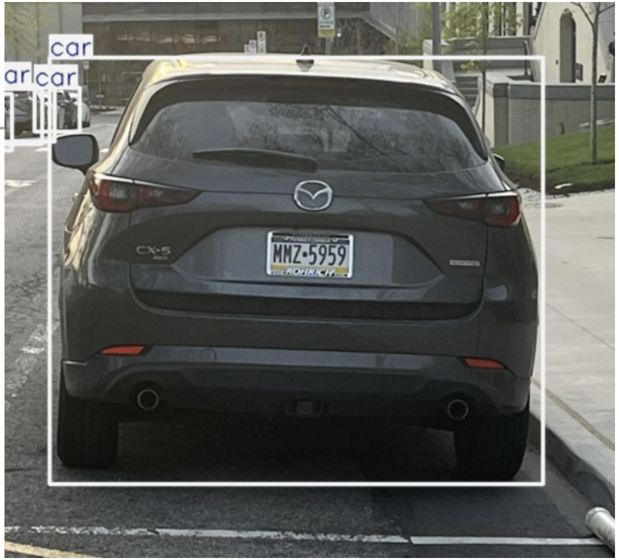

Testing, Verification, and Validation

Requirement	Testing Method	Testing Target	Result
Endurance and timing target	Run the raspberry pi for 55 minutes with our pipeline, continuously capturing images	• Every image is checked within 40 seconds • Power supply and cooling are adequate	• Every image was checked within 40 seconds • Power supply was adequate • Temp stayed below 75C
ML Model Pipeline precision and recall	Evaluate on images at different distances (5m, 10m) and weather/lighting conditions (day, night, rain, night+rain) of ideal image quality	• Final results meets 90% precision and 85% recall	• In progress • Initial testing shows 100% precision, 87% recall
GPS Accuracy	Get GPS location at 10 locations we know the true coordinates of outside (ex. CMU Flagpole)	• A result within 0.2 km for all	• All results were within 31 m

Testing, Verification, and Validation

Requirement	Testing Method	Testing Target	Result
Camera Quality	Evaluate through plate recognizer on images at different distances (5m, 10m) and weather/lighting conditions (day, night, rain, night+rain) taken with our camera	Meets 90% precision and 85% recall	In progress ● Many delays on getting camera ● Done some initial testing with our webcam
Edge-Cloud Communication	15 matches, should show up properly in supabase and 15 updates to the database should be added to the Edge database	Updates properly 100% of the time	Updated 100% of the time
System Test	Same as the ML Model Pipeline test, but with our camera's images	Meets 90% precision and 85% recall	In progress ● Waiting for images

Design Trade-Off Highlight

Trade-Off	Considerations/ requirements	Findings
Camera	• Infrared capability • High image quality • Autofocus to account for different distances • Enough FOV to read license plates from other lanes • Relatively low cost • Supports our edge device	• Raspberry Pi Camera module 3 supports all these features with a 12MP sensor at $25 • USB cameras we looked into didn't have high image quality and IR • Other raspberry pi cameras with higher resolution were significantly more expensive (ex. Arducam 64MP, $60, No IR) • Camera's with a zoom lens have lower FOV
Edge vs cloud ML inferencing	• Computing is cheaper and scalable in the cloud • Inferencing on cloud means sending images every 40 seconds; can't work without internet • Edge compute protects privacy	Hybrid Approach • Run faster, smaller models on edge for privacy and use without stable internet • Run larger model in cloud to verify the matches for performance • However, uses the most power

Design Trade-Off Highlight

Trade-Off	Considerations/ requirements	Findings
Edge device	• Has to be powerful enough to run our ML models • Must be compact enough to reasonably mount on a windshield • Must be able to run on car socket power	The Raspberry Pi 4 fulfills these criteria, on top of being cost efficient (5V 3A, $35). Other devices we looked of similar form factor at were at least $50. More powerful devices exist such as the Jetson Nano and Kria boards, they use more power, are harder to work with, and we do not use the performance gain.
Which OCR to use	• License plates may be from frontal, slanted, or rotated perspectives • There may be blurring due to weather conditions • Must work in different lighting conditions as well	Based on our testing, PaddleOCR performed better on images in the rain, snow, or blurred, and is the only one with built-in handling of slanting and rotation among the models we tested (PaddleOCR, EasyOCR, TesseractOCR).

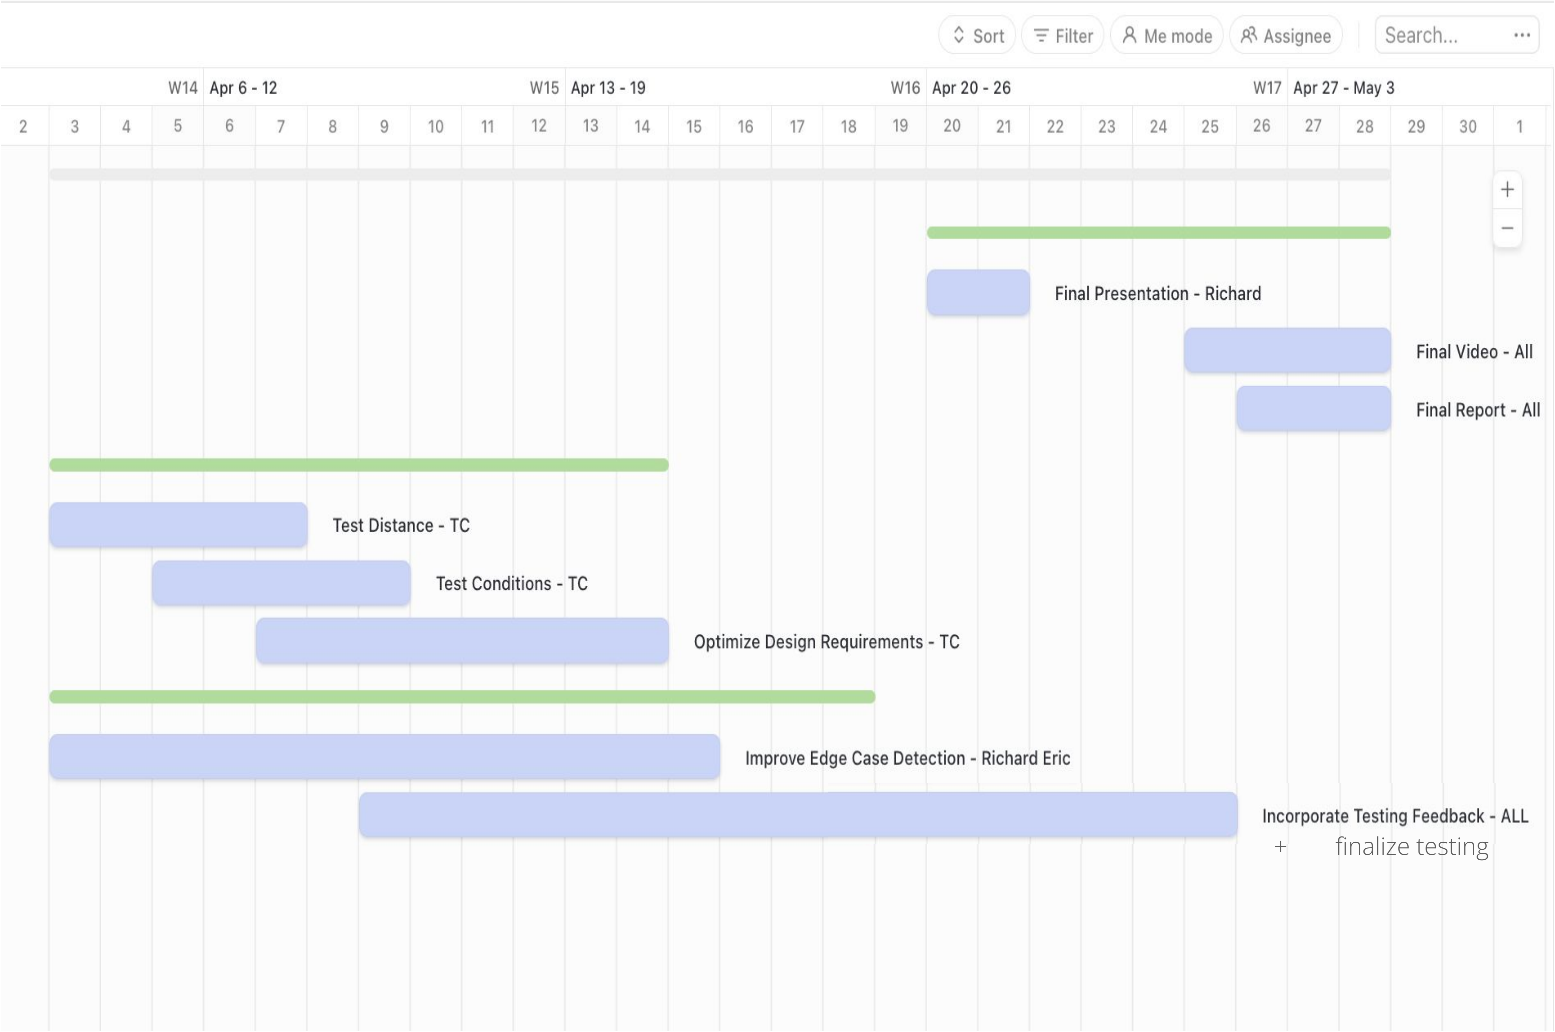

+ finalize testing