International Journal of Emerging Trends in Computer Science and Information Technology
ISSN: 3050-9246 | https://doi.org/10.63282/3050-9246.IJETCSIT-V7I2P105
Eureka Vision Publication | Volume 7, Issue 2, 27-37, 2026

Original Article

Predicting Speed-Limit Changes Using Smartphone Map Data before ADAS Camera Detection

Gaurav Pokharkar
Independent Researcher, USA.

Received On: 15/02/2026 ***Revised On: 21/03/2026*** ***Accepted On: 29/03/2026*** ***Published On: 06/04/2026***

Abstract - Advanced Driver Assistance Systems (ADAS) often rely on camera based traffic sign recognition to enforce speed limits, but vision alone can fail due to occlusion, adverse weather, poor visibility, or even missing signs. This paper proposes a complementary approach using smartphone based map data and enhanced localization (via Bluetooth and Wi-Fi) to predict upcoming speed limit transitions earlier and more reliably than onboard cameras. We present a comprehensive literature review of camera based speed sign recognition limitations, analyze the coverage and accuracy of digital speed limit maps (Google, Apple, HERE, OpenStreetMap), and evaluate how Bluetooth Low Energy (BLE) beacons and Wi-Fi fingerprinting can refine vehicle positioning in challenging environments. We design a predictive model that computes time to speed limit change using map segment metadata, current vehicle dynamics, and localization uncertainty. A conflict detection and resolution framework is outlined to cross verify map predictions with camera readings, assign confidence scores, and correct errors in real time. Finally, we propose an experimental field evaluation across multiple mapping providers (urban, rural, highway, tunnels, and con- struction zones) to benchmark prediction accuracy, lead time gains, localization effects, and failure modes. The results indicate that integrating map based speed limit data with enhanced localization and sensor fusion can significantly improve speed limit awareness, providing earlier warnings (several seconds ahead of camera detection) and greater reliability in diverse conditions. We discuss system architecture, algorithm pseudo code, a comparative performance table of mapping platforms, and consider practical deployment issues (data freshness, con- nectivity, privacy, and regulatory compliance).

Keywords - ADAS, Autonomous Vehicles, Localization, Map-Ping, Intelligent Adaptive Cruise Contorl (Iacc).

1. Introduction

Camera based Traffic Sign Recognition (TSR) systems have matured with machine learning (e.g. YOLO, SVM classifiers) and achieve high accuracy in ideal conditions [1]. However, their performance deteriorates in non-ideal conditions com- monly encountered on roads. Occlusion is a major issue: a speed limit sign partially or fully blocked by a truck, bus, or foliage can go undetected [2] . Even moderate occlusions can cause a drastic drop in detection rates (e.g. a 41% drop in mean average precision under moderate camera occlusion was reported in one study) [3]. Adverse weather and lighting also impede recognition. Rain, fog, or snow not only obscure the camera’s view but introduce noise (blur, glare, low contrast) that confuses image based algorithms [4], [5]. For instance, a recent evaluation found that older style reflective signs were often missed by a vision system at night due to poor reflectivity [5]. Similarly, heavy rain or snow can reduce camera detection ranges and accuracy dramatically, as sensors lose clarity.

False positives and negatives plague camera only systems in complex environments. A common false positive scenario is when the camera ”sees” a speed number that isn’t an actual regulatory sign for example, an advertising billboard or a speed sticker on a vehicle can mimic a speed limit sign. An example from prior research is shown in Figure 1, where a temporary sign with an arrow was mis-read as a standalone speed limit; more advanced processing to detect the arrow was needed to avoid a false alert [6]. Cameras can also misread adjacent road signs e.g. a highway off ramp’s 40 mph sign might be picked up by a car on the main road, causing a false speed warning. Conversely, false negatives occur when legitimate signs are missed: e.g. a camera might fail to detect a sign that is above the driver’s eyeline or one that’s faded or dirty. Glare from low sun angles or headlight reflections at night can wash out sign numbers, leading the system to miss the sign altogether [4]. In construction zones or where temporary electronic speed signs are used, camera algorithms trained on standard sign shapes may not recognize LED numeric displays or unusual sign placements, resulting in missed detections. Edge cases also include conditional speed limits (time of day, weather based). A camera might read the sign’s numeric value but not contextual text (e.g. ”School days 7–9am”) and thus could present a misleading limit. In summary, vision only speed limit recognition suffers in many real world edge cases: occlusion, poor weather/lighting, nonstandard or missing signage, and contextual complexity all contribute to degraded reliability [4], [7]. These limitations motivate using an alternative source of truth the digital map to fill in or predict speed limits when the camera cannot

2. Feasibility of Map API Based Predictive Speed Limit Awareness

Digital mapping platforms today maintain extensive road databases, including attributes like speed limits. Tapping into these via smartphone APIs could allow a vehicle to know the posted speed ahead even before a sign comes into view. However, the usefulness of this approach hinges on the com- pleteness, freshness, and accuracy of the map providers' speed limit data. We examine four major sources:

2.1. Google Maps Platform (Roads API)

Google's Roads API can return the speed limit for a given road segment when provided with GPS coordinates or a path [8]. Google's speed limit coverage has expanded globally to 47 countries (as of 2025), heavily focused on North America, Europe, and developed regions [9]. Notably, only 8 of those 47 are in the "global south," indicating large gaps in developing countries [9]. Google employs a mix of sources to populate speed data: official datasets from local governments (where available) and machine vision from Street View imagery processed by AI [10]. The AI approach attempts to read speed limit signs from street level images worldwide, even handling different sign styles and conditional limits [10]. This helps keep data current Google reports using imagery plus analysis of traffic speeds to infer changes (e.g. if drivers consistently slow down on a segment, the system flags a possible new lower limit) [10]. Despite these efforts, Google acknowledges varying accuracy: speed limits come from "several sources that vary in accuracy and coverage" [10]. Quality issues range from outright missing data on smaller roads to incorrect values; Google allows users to report speed limit errors via the Maps app for correction [10]. Overall, Google's speed limit dataset is among the most extensive, but not infallible errors in certain areas have been noted (e.g. early versions of Google's data would sometimes show default 25 mph in places that were actually 35 mph zones). Continuous AI driven updates aim to reach the 90% accuracy required by emerging European ISA legislation, but achieving this solely by vision is challenging [11], hence Google's emphasis on fusion of data sources

2.2. Apple Maps

Apple's mapping platform similarly provides speed limit alerts in its Maps app and via Apple CarPlay, though program- matic access is less open than Google's. According to Apple's iOS feature availability, "Maps: Speed Limits" are available in 40 countries (only 1 of which is a developing nation) [9], [12] . Apple likely obtains speed limits from TomTom and local authorities, as well as its Apple Maps image fleet. The data coverage is broad across Europe and the Americas (e.g. Andorra to the United States are covered) [12]. Users of Apple Maps and Waze (which is Google owned) have observed that Apple's speed limits sometimes differ from Google's on the same road, pointing to differences in data sources or update frequency. The World Bank analysis highlighted that Apple's coverage, while good in many regions, had virtually no coverage in many low income countries [9]. Apple does not publicly detail its update processes, but given their fewer on the ground resources compared to Google, some data may lag. Anecdotally, Apple Maps was slower to roll out speed limits (which it added around iOS 14 in 2020) and tends to focus on major roads first. For our purposes, Apple's data is a valuable alternative in iOS centric vehicles, but one must be mindful of potential gaps outside the supported regions.

2.3. HERE Maps SDK

HERE Technologies (formerly Navteq/Nokia) is a primary supplier of automotive grade map data, including speed limits, to many OEM navigation systems. HERE's data is known for being highly structured and audited their map platform covers 100+ countries with speed limits for nearly all classified roads. They combine official road authority data with their own fleet surveys. One study by the World Bank team in Latin America found that HERE (and TomTom) had far more complete speed coverage for city streets than community sourced maps, but even HERE's data sometimes disagreed with reality by significant margins [9]. For example, in Mexico City, officials noted cases where HERE's map listed a street at 80 km/h while the actual posted limit was 50 km/h [9]. These discrepancies can arise if a speed limit was recently changed by local ordinance and not updated yet in the map. Nonetheless, HERE's strength is its automotive focus its data powers ISA systems in many European cars and has update mechanisms (including over the air updates to vehicles) to remain compliant with regulations. European NCAP ratings favor map informed ISA and require that digital maps used for speed limits be updated at least quarterly [11]. HERE and similar providers achieve this via subscriptions. Therefore, using HERE SDK in a smartphone (through an app) could provide a robust speed limit feed. The feasibility is high if the vehicle or device has access to HERE's latest map; however, HERE data is proprietary and typically requires licensing, which might limit its use in a purely crowdsourced smartphone scenario.

2.4. OpenStreetMap (OSM)

The open source mapping community offers speed limit data via the maxspeed tag on road segments. OSM's advantage is global coverage and openness; any user can contribute updated speed info, which makes it very fresh in some lo- cales (volunteers often update changes within days). However, completeness is inconsistent. In well mapped cities of Europe or North America, a large proportion of roads have maxspeed tags, but elsewhere the coverage can be sparse [13]. A study in Phoenix, AZ found that in OSM, a whopping 47% of road segments had no speed limit tagged at all [13]. Our own analysis in Latin American capitals revealed many residential streets missing explicit speed values (default limits apply, but those defaults may not be recorded in OSM). Figure 1 below (from a World Bank study) dramatically illustrates the platform inconsistency: two different navigation apps (using different data sources) show 20 km/h vs 80 km/h as the "speed limit" for the same avenue, neither matching the actual posted 50 km/h sign [9]. Such disparities often involve OSM or outdated datasets. Quality control is another concern OSM relies on consensus, so errors (e.g. a user accidentally tagging a highway as 50 mph instead of 50 km/h) can persist until noticed. On the other hand, OSM has rich metadata capabilities: it can encode conditional limits (time based, weather based) through tags like ”maxspeed:conditional” (e.g. ”40 mph @ (school_hours)”). If utilized, this could convey nuance that standard APIs might not. In practice though, conditional tags are underused and not widely parsed by consumer apps yet. Freshness of OSM is generally excellent where active contributors exist e.g. in some European cities, volunteer mappers update speed zones soon after new laws. But without a guaranteed update schedule, relying solely on OSM could miss recent changes elsewhere. For ISA use, OSM can be a valuable reference (especially as an open data source to compare against proprietary data), but one must account for its gaps. Efforts are underway (like the World Bank’s pilot in Latin America) to help cities digitize official speed limits and feed them into OSM or open databases [9], which will improve completeness.

Figure 1. False Positive Speed Limit Detection [6]

In summary, using smartphone accessible map data for predictive speed advice is feasible but demands careful integration of multiple sources to overcome each source’s limitations. Google and Apple provide easy to use APIs but have patchy coverage in some regions. HERE offers high accuracy but requires licensing. OSM is free and often up to date, but not uniformly complete. A prudent strategy is to fuse these:

E.g. use Google’s API as primary (for its convenience and global usage in apps), cross check against OSM data (to flag discrepancies or fill blanks), and wherever possible leverage authoritative sources (HERE or official city data) for critical segments. Ultimately, map based speed limit info can signif- icantly augment camera based systems in fact, recent Euro NCAP tests show that vehicles using map data for ISA scored higher on speed assistance performance than those using camera alone [11]. Map data particularly shines in handling implicit or conditional limits that cameras often fail to interpret (e.g. unposted urban default speeds, school zone times) tests found that only 20% of camera only systems correctly handled time based or implicit limits, whereas map informed systems can easily include those rules [11]. However, to truly trust map data, one needs confidence in its accuracy; thus a theme of this research is not only using maps but also continuously validating and updating them using vehicle sensor feedback.

3. Enhanced Localization with Bluetooth and WI-FI

A key requirement for predicting a speed limit change ahead is knowing precisely where the vehicle is on the map. In open sky conditions, GPS/GNSS typically provides 3–5 m accuracy which is sufficient to determine if you are, say, 50 m from a speed zone boundary. But in urban canyons (downtown high rises), underpasses, tunnels, or dense forests, GNSS accuracy degrades (10–50 m errors or complete drop outs). Relying on GPS alone could make the timing of speed change alerts very unreliable e.g. a 30 m position error at 100 km/h means a ±1 second timing uncertainty. Bluetooth Low Energy (BLE) beacon and Wi-Fi fingerprinting techniques can augment localization in these scenarios by leveraging local signal signatures independent of GPS [14].

3.1. Bluetooth Low Energy (BLE)

BLE beacons are low power transmitters (often used in malls, etc.) that broadcast identifiers. By planting BLE beacons along a road (or leveraging existing ones in infrastructure), a vehicle’s smartphone can pick up their signals. The received signal strength indicator (RSSI) can provide a rough range estimate to the beacon. With multiple beacons in range, one can trilaterate a more precise position. Even with a single beacon, some advanced methods use the antenna array on a modern smartphone to estimate angle of arrival (DOA) of the BLE signal [15]. Research has shown that BLE based localization can reach accuracies of a few meters in indoor settings and could be similar outdoors given known beacon positions [14]. For example, one 2024 study achieved room level (within 3 m) accuracy using BLE fingerprints in an office [16], and other experiments report sub 5 m accuracy for BLE positioning with proper calibration [14]. In a vehicle context, a practical deployment might include BLE beacons at tunnel entrances, along long bridge spans, or in urban downtown cores places where GPS falters. The smartphone could contain a database of beacon IDs and their coordinates (similar to how Assisted GPS servers have databases of cell tower locations). When it detects certain BLE IDs and signal strengths, it can snap or correct its position on the map. BLE has advantages: the beacons are cheap and can be deployed by road authorities or even crowd sourced (e.g. businesses along a road could have beacons). It also works underground (e.g. in parking structures or tunnels where radio can propagate but GPS doesn’t). The downside is the need for infrastructure and maintenance (battery powered beacons need replacement or power supply). Also, RSSI distance estimation is notoriously noisy due to interference and multi path, especially with a metal vehicle body around filtering and calibration are needed to use BLE reliably. Newer approaches use BLE angle of arrival with multi antenna receivers to greatly improve accuracy [15], which could be leveraged in vehicles with multi antenna Bluetooth setups.

3.2. Wi-Fi Fingerprinting

Nearly every smartphone continuously scans for Wi-Fi net- works. In urban areas, the density of Wi-Fi access points (APs) is high, and their signals can serve as location "fingerprints." Wi-Fi positioning typically works by comparing the list of visible networks and their signal strengths to a radio map (a database of what Wi-Fi signals are expected where) [17]. Companies like Google and Apple already use Wi-Fi signatures for location services (that's how your phone finds location quickly via "Wi-Fi & mobile network location" even before GPS locks). For our application, Wi-Fi could significantly enhance localization in city streets. For example, if GPS is bouncing among tall buildings, the phone's Wi-Fi scan might reveal, say, the MAC address of a cafe's router which a cloud database knows is at a fixed location thus correcting the position. Studies show that urban Wi-Fi positioning can achieve 5–10 m median accuracy outdoors [14], and as fine as 2–3 m in optimized cases or indoor settings [14]. One experimental system replicated GPS like behavior in a city by using a war driving collected Wi-Fi map and achieved 2.25 m mean error in open areas, and about 0.5 m when fused with inertial sensors and particle filters [14]. These results are promising for lane level accuracy. The feasibility of using Wi-Fi signals is high because no extra hardware is needed (just the phone's Wi-Fi radio), and it works wherever people install Wi-Fi (virtually everywhere in cities, and even along highways one might catch signals from nearby buildings). The challenges include needing an up to date radio map: Wi-Fi networks can appear, disappear, or change IDs (SSID, BSSID) over time. However, crowdsourcing can address this e.g. Google continuously updates its Wi-Fi location database via devices running Google Maps. For our ISA system, we can leverage those existing databases through APIs (Android and iOS location services abstractly use Wi-Fi/cell data to improve GPS). Alternatively, a dedicated fingerprint approach could be used in particularly troublesome spots: for instance, mapping the Wi-Fi signature in a long tunnel and programming the system to know "if you see these 3 Wi-Fi SSIDs strongly, you are 200 m inside the tunnel, so likely halfway to the tunnel's speed zone end."

3.3. Integration and Complementarity

Bluetooth and Wi-Fi localization are complementary to GNSS. BLE beacons can give absolute references at specific points (like "you are at tunnel marker 3"). Wi-Fi gives a continuous position estimate wherever networks are present. Both can be fused with GNSS and inertial measurement unit (IMU) data using sensor fusion algorithms (e.g. an Extended Kalman Filter or particle filter) [14]. The result is a more robust positioning module for the vehicle. This improved position directly translates to better prediction of when a mapped speed limit change is coming. Instead of the car guessing its distance to a speed zone based on a flaky GPS fix, it will "know" it with higher confidence even in urban downtown, it might localize to the correct block or lane. Another benefit: position confidence estimation the system canoutput an uncertainty (e.g. ±5 m). Our prediction model (next section) will use that to adjust how early to alert a change (larger uncertainty ¿ more conservative early warning).

It's worth noting potential limitations: interference and scalability. In a dense city, using Wi-Fi/BLE means dealing with interference from many devices. As vehicles adopt V2X communications on similar bands, the 2.4 GHz spectrum will be noisy. Moreover, if many vehicles simultaneously scan and query location services, there could be network or privacy concerns. Studies caution that Wi-Fi/BLE accuracy can drop when too many devices or reflections are around [14]. Nonetheless, modern approaches (e.g. Wi-Fi Round Trip Time ranging (IEEE 802.11mc) which some call " Wi-Fi FTM") can even provide direct distance measurements to APs with sub meter accuracy and are starting to appear in phones [18]. Such techniques, combined with conventional fingerprinting, promise lane level localization in GPS denied environments [18], [19] a crucial enabler for map based speed predictions to be trustworthy. We will assume our system can obtain a location with 5 m or better accuracy most of the time by fusing GPS, IMU, BLE, and Wi-Fi inputs.

4. System Architecture for Predictive Speed

4.1. Limit Assistance

To combine all these elements, we propose an integrated system architecture (Figure 2) for map augmented speed limit awareness. The architecture comprises several modules working in a closed loop:

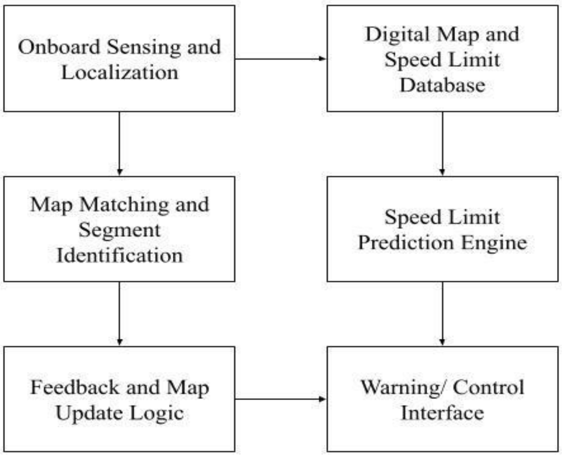

Figure 2. System Architecture Diagram

- Onboard Sensing & Localization: The vehicle (or driver's smartphone) continuously gathers positioning data using GNSS, IMU, Wi-Fi, BLE, etc. A sensor fusion algorithm (e.g. Kalman filter) produces a real time vehicle position and heading on the digital map, along with an estimation uncertainty (e.g. covariance) [14]. Simultaneously, an onboard camera system detects and recognizes speed limit signs in the vehicle's forward view when visible. Each detection includes the sign's content (e.g. "45 MPH") and confidence level, and approximate range ahead.
- Digital Map & Database: This module contains the road network graph with associated speed limit metadata. It can query multiple sources for instance, primary query to Google's Roads API for the current road, fallback to HERE/OSM data if needed. This database isn't static; it can be updated OTA (over the air) or from cloud services. The map

also encodes conditional rules (time based limits, average speed zones, etc.) where available.

- Map Matching and Segment Identification: Using the precise location from module 1, the system identifies which road segment the vehicle is on and finds upcoming segments. A map matching algorithm (often a Hidden Markov Model using sequential GPS points) corrects for any off road GPS errors and locks the vehicle to the correct road in the map [6]. Knowing the route ahead (if navigation is active) or simply following the road graph if not, the system looks ahead along the path to find the next segment where the speed limit changes. For example, if you are in a 65 mph zone and 500 m ahead it becomes 55 mph, the map provides that distance.
- Speed Limit Prediction Engine: This core module takes the upcoming change information (distance Δd to the next speed zone and the new speed value) and computes a Time to Speed Change (TTSC) prediction. The algorithm accounts for the current vehicle speed v and acceleration a (from OBD or phone sensors) if accelerating, the remaining distance is covered more quickly; if braking or climbing a hill, more slowly. It also incorporates position uncertainty: if the system is uncertain by δ, it assumes the shorter effective distance for safety. The output may be a prediction such as “Speed limit drops to 50 km/h in 15 seconds,” or an equivalent anticipation distance. A simplified pseudo code example for this prediction is provided below. The model can also adapt: if the vehicle is traveling well below the limit (for example, due to traffic), the urgency decreases, whereas if it is at or above the limit, the system may recommend easing off the throttle earlier.
- Warning/Control Interface: Using the prediction from module 4, the system can do one of two things in an ISA context: (a) Driver Alert e.g. display “Upcoming 50 km/h zone” with a countdown, and possibly an auditory chime if the driver is still above that speed when entering the zone. Or (b) Active Control if integrated with adaptive cruise or speed limiting, the system could smoothly decelerate the vehicle before reaching the new zone (similar to how some modern cars use map data to adjust cruise setpoints). The interface also shows the current speed limit and the next one (for driver’ssituational awareness).
- Feedback & Correction: Here lies an important differ- entiator of our syste it doesn’t passively trust the map, it actively validates and corrects it. When the camera (from module 1) detects a speed sign, this information is sent to a Data Fusion & Conflict Resolver component. If the camera’s recognized limit differs from the map’s expected value for the current location, the system flags a discrepancy. A confidence logic (described in the next section) decides which source to believe or whether additional evidence is needed. High confidence camera detections can trigger an immediate override of the map data (e.g. “Map said 80, but camera clearly saw 60, cor- rect to 60”). The discrepancy can be logged and reported to a cloud service for map improvement (crowdsourced update). This feedback loop allows the digital map to learn from real world ground truth, steadily increasing its reliability [20]. Similarly, if no sign is seen where the map expected one, the system could verify if the sign might be missing physically and adjust accordingly (for example, if three cars all pass where a “50” sign is mapped and none see it, perhaps it was removed and the map is outdated).

The overall architecture thus merges on board perception with off board knowledge, exploiting the strengths of each. As noted in prior ADAS research, combining vehicle sensors and infrastructure data yields the most robust performance [6], [21]. The camera provides immediacy and real time validation, while the map provides foresight behind visual range and coverage when vision fails. In the next sections, we detail the predictive algorithm and the map data conflict resolution mechanism.

4.2. Pseudo code for Predicting Time to Speed Change

This simplified model predicts the time remaining to a speed limit change. In practice, one might incorporate speed profiles: e.g., if the upcoming limit is lower (Ln ¡ Lc), many ISA systems would start to gently decelerate perhaps 5–10 seconds before reaching the sign [11]. Our model can be extended to recommend a deceleration trajectory: using current speed and required speed drop, compute a comfortable deceleration (e.g. 1–2 m/s²) and distance needed. If the available distance is shorter, the system could alert the driver earlier or more urgently. Conversely, if the limit will increase (say from 50 to 80 km/h ahead), the system might not need an alert but could preemptively allow acceleration after the sign. Timing also ties into localization uncertainty as shown, we subtract a margin so that if we might be closer to the sign than GPS thinks, we err on the side of caution (earlier alert).

Finally, all these computations run continuously (for exam- ple, updating t_{pred} every second). If the vehicle takes a turn or follows a new route, the map matching module updates Δd and L_n accordingly. The architecture is designed to be modular and real time, suitable for deployment in modern vehicles or as a smartphone application paired with an OBD II reader for speed data.

Algorithm 1 Time to Speed Change Prediction

```
Require: v (current speed), a (acceleration), $\Delta d$ (distance to change), $L_c$ (current limit), $L_n$ (upcoming limit), $\sigma_{pos}$ (localization std. dev.)
$d_{eff} \leftarrow \Delta d - 2\sigma_{pos}$
$d_{eff}$   max($d_{eff}$, 0)
if $\Delta d$   0 or $L_n$ = None then
    return None                ▷ no prediction needed
end if
```

```
if $|a| < 0.1$ then
    $t_{pred} \leftarrow d_{eff}/\max(v, 0.1)$
else
    $discr \leftarrow v^2 + 2ad_{eff}$
    if $discr < 0$ then
        $t_{pred} \leftarrow d_{eff}/\max(v, 0.1)$
    else
        $t_{pred} \leftarrow \frac{-v + \sqrt{discr}}{a}$
        if $t_{pred} < 0$ then
            $t_{pred} \leftarrow \frac{-v - \sqrt{discr}}{a}$
        end if
    end if
end if
$t_{pred} \leftarrow \max(t_{pred}, 0)$
$t_{pred} \leftarrow t_{pred} + 1.0$    ▷ latency / reaction buffer
return $t_{pred}$
```

Stances. Confidence factors include: (1) Camera confidence the recognition algorithm typically provides a confidence score for the detected sign and classification. If the sign was fully visible, close, and the system is, say, 99% sure it read “60”, that’s strong evidence. If it was a partial/angled view, confidence might be lower. (2) Map confidence not all map data is equal. The system can store a quality metric for map info, for example: data obtained from official updates within last 3 months is high confidence, whereas data inferred by AI or crowd sourced could be labeled lower confidence. Also, if other vehicles or the cloud have previously flagged this segment as problematic, we downgrade confidence. (3) Contextual cues if the sign detected by camera has modifiers (like an arrow, or time plaque), the map might not encode that nuance, so the conflict might be only partial (map says 60 always, camera sees 60 with “7 9AM” which currently doesn’t apply; technically not a conflict for current time). The system should recognize conditional signs and handle them by time rules rather than treating as a direct conflict. Another cue: road type and location. If the map says 80 km/h on what visibly is an urban street, one might inherently suspect the map is wrong (80 in city center is unlikely) [9].

Using these factors, the system might assign a probability that the true speed limit is the camera’s value vs the map’s. For instance: P(true=60—camera) = 0.99 if the sign was clear, while P(true=80—map) = 0.5 if the map data is older or contradicts usual speed patterns. There is also the possibility both are correct in some way e.g. a hidden sign indicated

5. Detecting and Correcting Map Errors in Real Time

One of the most significant challenges in relying on digital maps for speed limits is map errors or outdated information. Our system addresses this by actively comparing map predictions with ground truth from the camera and other sensors, and then applying logic to reconcile differences. This process involves conflict detection, confidence scoring, and if needed real time correction of the speed limit used by the ISA system.

Conflict Detection: Whenever the vehicle’s camera detects a speed limit sign, the recognized value (say X km/h) is compared to the map’s current believed speed limit for that road segment. A conflict is flagged if there is a mismatch (e.g. map says 80, sign says 60). Additionally, the absence of a sign where the map expects one can be a conflict: for instance, map data might indicate a change to 30 mph should have occurred (perhaps a city boundary) but the camera never saw a 30 sign. This could mean the map’s data is wrong (no such change exists or the sign was removed). Other sensors or logic help here: if the vehicle traveled the entire distance of the supposed speed zone without seeing a sign, likely the map is wrong. Another conflict scenario is multiple signs: camera might see an “advisory” speed sign (yellow in US, not legal limit) which map might not have the system should distinguish regulatory signs from warnings.

Confidence Scoring: Once a conflict is detected, the system evaluates which source is more likely correct under the circum- a different limit on an intersecting road, which the camera picked up erroneously. But if the map matching is correct, that shouldn’t happen often (the system knows which road we’re on).

Resolution and Correction: Once the system leans one way, it will act accordingly. If the camera is deemed correct (common case: new lower speed limit not yet in map), the system will override the map’s value for all further calculations (the driver will be alerted based on the camera detected 60 km/h, not the outdated 80). In parallel, it logs this discrepancy. Many modern cars do this internally for example, VW’s traffic sign recognition can supplement the nav’s data and display the latest seen sign if it differs. Our system goes further by planning to update the map: it can send the detected change to a cloud service or mark it for user verification. For an open source approach, it could even format an OpenStreetMap change suggestion (e.g. “maxspeed on Road X should be 60, not 80”) to be reviewed by the community. For proprietary maps, these reports could be channeled to the provider (Google Map’s user reports as described, or HERE’s feedback loop). Over time, this ensures the digital map is corrected essential since studies have found 3–6% error rates in official speed limit maps even in developed cities [20], often due to changes not yet reflected. As reported by Nexar’s 2022 Phoenix study, about 2.5% of speed limits change per year in typical cities [20]. So without automated updates, a map could be 97.5% accurate after a year, which may be insufficient for high safety requirements [20]. Our system’s feedback aims to push that closer to 100% by catching changes soon after they occur (crowdsourced from drivers).

If the map is deemed correct and the camera likely false (say camera read a truck’s speed sticker or a sign from another road), then the system will ignore the camera reading in terms of setting the limit. However, it might still display it to the driver momentarily (some vehicles flash a sign they saw with a question mark if unsure). The system could also cross check: if one camera reading conflicts but shortly after another sign is seen confirming the map, it reinforces that the first was false. Use of multiple sensor modalities can help too e.g. if the vehicle has a database of common false sign triggers (like “end of speed limit” signs or truck speed markings), it can classify those. Furthermore, vehicle to vehicle (V2V) communication could play a role: if other cars ahead recently confirmed the map’s limit, a lone false camera hit can be discounted. This is beyond our current scope but worth mentioning as future enhancement (crowd validation in real time).

Continuous Map Reliability Evaluation: Even with no im- mediate conflicts, the system monitors over time how often the map has been “right” vs “wrong” in predicting upcoming speed zones. If certain areas have frequent corrections, it might treat the whole region’s map data as suspect until updated. In academic research, reliability scores or trust models are created for map data [6], [20]. For example, an algorithm could assign a confidence level to each map segment’s speed limit attribute, updating that confidence when a vehicle’s observation agrees or disagrees. An average agreement rate can be computed. If a segment consistently matches camera observations (high confidence), the system might choose to alert a driver slightly later, trusting the map timing. If confidence is low, it will err on earlier, more cautious alerts. In essence, the map stops being a static input and becomes a dynamic data layer that the vehicle helps to validate.

There is also the scenario of temporary speed limits e.g. work zones. These often are not in any map database because they are short term (maybe a 2 week construction). A camera might catch a “40 km/h” orange sign in a construction zone whereas the map says 60. Here, the map is not “wrong” per se, it’s just unaware. Our conflict resolution should treat this as a legitimate new limit (at least for the time being). The system should immediately trust the camera in this case (since maps rarely include temporary limits) and possibly engage a mode to periodically check if the construction zone is over (perhaps expect an “End road work” sign or use geofenced data from public feeds). For updating maps, this could be tricky some mapping companies are exploring real time work zone feeds from city DOTs to incorporate into navigation. If available, that could be another input to our map data to avoid treating it as an error.

In summary, map data error handling in our system uses sensor fusion and logic to ensure the vehicle always follows the correct speed limit, whether the information came from digital map or physical sign. The strategy is fail safe: if there’s doubt, choose the lower speed limit to avoid speeding. By continuously learning from discrepancies, the map data improves, enabling even better prediction for the next driver. This synergy between onboard vision and cloud based map updates is a powerful feedback loop envisioned by many researchers to keep HD maps up to date [20]. It also provides a path to eventually predicting even temporary changes if many users report a new work zone speed, a predictive system could warn upcoming drivers of it before their cameras even see the first sign or cones.

6. Experimental Evaluation across Mapping Providers

To validate the proposed system, a comprehensive field test regime is designed. The evaluation aims to compare prediction accuracy and timeliness against camera only detection, across different mapping data sources, in multiple road environments. Test Scenarios: We identify several representative driving environments to test: (a) Urban streets with frequent speed limit changes (e.g. 50 km/h to 30 km/h near schools, etc.), tall buildings (to challenge GPS), and potentially missing signs (some residential areas rely on default speed, one sign at neighborhood entry). (b) Rural highways long segments with infrequent signs, occasional small towns where the limit drops (e.g. 90 to 50), and maybe ambiguous signage (signs can be far apart). (c) Motorways/Highways high speed roads, overhead or roadside signs, including cases of variable speed limits (electronic signs) and ramps (to test if the system doesn’t confuse adjacent roads). (d) Tunnels and Overpasses segments where GPS is lost and reliance on inertial/ Wi-Fi is tested; typically speed drops in tunnels (e.g. 100 outside, 80 in tunnel). (e) Construction zones temporary speed limit areas, to test detection of map vs reality mismatch. We will also include night and inclement weather drives to test the robustness when the camera might fail more often and we lean on the map.

Procedure: A test vehicle is equipped with the integrated system (or a data logging setup with all components). For each scenario, the vehicle is driven through a course where known speed limit changes occur. The system’s behavior is logged: at what distance/time before the actual sign did it alert or take action? What speed limit did it predict versus what was observed? We do this for multiple mapping provider configurations: e.g. run the system using Google’s map data, then using OSM data, then HERE data, etc., to see differences. The vehicle’s camera detection alone (no map aid) serves as a baseline essentially, the point at which the camera definitively reads the sign is “time zero” for purely vision based ISA. We measure the lead time our predictive system achieves (how many seconds earlier it knew the new speed) compared to that baseline.

Metrics Collected:

- Prediction Lead Time: For each speed transition, the time (or distance) before reaching the sign at which the system first indicated the upcoming new limit. For example, if approaching a 60 40 drop, the camera may detect the 40 sign at 50 m (about 2 s before passing it), while the map based system may detect it 200 m (8 s) earlier. We record the improvement (6 s in this case) and average this across scenarios.
- Accuracy of Predicted Limit: Whether the system’s predicted upcoming limit was correct. If the map indicates the next zone is 40 but the true limit is 50

(or unchanged), the prediction is incorrect. Such errors reflect map data quality — well maintained maps should exhibit few errors; incomplete maps will show more.

- Frequency of Missed Changes: A missed change occurs when the system fails to predict a real speed limit change (for example, a temporary 30 km/h construction zone missing from the map and unseen by the camera due to occlusion). False alarms (predicting a change that does not occur) are also counted.
- Localization Error Impact: Tests include intentional perturbations to localization (e.g., offsetting GPS by
- +30 m) to examine how timing errors arise. Comparing actual distance to sign with the system's assumed distance at alert time reveals the timing error caused by localiza- tion uncertainty. With Wi-Fi/BLE augmentation, timing error should ideally remain below 1 s.
- Provider Comparison: Each mapping provider is eval- uated across the above metrics. For instance, OSM may miss more transitions due to incomplete tagging; HERE may have very few errors but slightly shorter lead time; Google may perform strongly in regions with dense Street View but weaker elsewhere. Example results might in- clude: Google — 95% correct, 5.0 s average lead; HERE 98% correct, 4.5 s lead; OSM — 85% correct, 5.5 s lead. A comparative performance table will summarize these findings.

Benchmarking and Baselines: We use camera only detection as one baseline. Another baseline could be a fixed distance warning approach (some ISA just beep when you cross a sign). Our system should outperform both by giving earlier and more reliable cues. We also compare to ground truth: using a high precision GPS (RTK) and manual observation of signs to know exactly when a change happened. This ground truth allows calculation of any early/late warnings precisely. If the system warns 3 s early on average with ±1 s variance, that might be acceptable; if sometimes it warns too late (after passing the sign), that's a failure.

We also incorporate failure case analysis. Each time the system failed (miss or false alarm), we diagnose why. Was it because the map was wrong or outdated? Or because localization was off and it triggered late? Or perhaps the camera was needed but it was blinded by heavy rain? By categorizing failures, we can refine the system. For example, if many failures in construction zones come from missing temporary limits, we might suggest integrating official work zone data feeds. If tunnels caused a late alert due to losing GPS for a bit, we might improve the inertial dead reckoning or place BLE beacons as suggested.

- Multiple Provider Field Results Hypothetical Example: In an urban test route with ten known speed limit transitions, Provider A (e.g., Google) correctly predicted nine of them, with an average alert distance of approximately 120 m (about 5 s) before the sign. One transition was missed because the map still contained an outdated value (the system warned for 50 km/h while the posted sign was 40 km/h, corrected only after the camera observed it).

Using Provider B (OSM) on the same route, seven transitions were predicted correctly. Several residential streets lacked mapped speed limits, resulting in no warnings for those drops; the camera handled them instead. However, OSM contained one correct speed zone that Google missed (likely due to a recent local mapping update), illustrating the benefit of combining multiple data sources.

Provider C (HERE) might achieve 10/10 correct predictions (all changes mapped), though alerts may occur slightly later for example, around 100 m before the sign possibly due to conservative segmentation of speed limit boundaries in their dataset.

These differences would be summarized in a comparison table, with notes such as: "OSM failed on three changes due to missing tags; HERE mis timed two events, likely from a 20 m map offset; Google had one outdated entry." Such observations help determine which provider, or combination of providers, offers the most reliable overall performance.

Lead Time vs. Camera Detection: One key outcome we expect is that map based predictions give a significantly longer lead time than camera alone. If a typical camera detects a sign at 50 m (depending on sign size, camera resolution), at 50 km/h that's about 3.6 s before reaching it. Our map could know of the change even a few hundred meters ahead (limited by how far ahead we look we could look 1–2 km ahead if needed, but practically we might focus on the next change). So we might be giving 10–15 s notice on highways, and maybe 5–8 s on urban streets (because changes are closer together). This extra reaction time can be valuable for driver comfort and safety no more last second braking when a sign appears. We'll quantify this benefit.

Localization Benefits: In our tunnel tests, for instance, we expect that using Wi-Fi fingerprinting, the system still correctly transitions to the tunnel's lower speed limit at the right point, whereas a GPS only approach might either lose position or assume the change at the wrong time. This can be shown by plotting the vehicle's perceived position vs actual. With augmentation, the error is small, without it, the position might jump, causing a late or early alert. Quantitatively, with augmentation, maybe the average speed change timing error is ¡0.5 s; without, it was 2 s. This demonstrates the necessity of the localization module.

Driver Response and Human Factors: Although mainly a technical evaluation, we should note if early warnings actually translated to smoother driving. In tests where a human driver is reacting, does the extra lead time result in gentler decel- eration? We can monitor vehicle deceleration profiles: ideally map informed ISA leads to earlier, less abrupt slowdowns, whereas camera only might cause harder braking when the sign “pops up.” If we had multiple drivers, we could even gather subjective feedback did they feel more confident with the system on? This is more anecdotal but important for real world acceptance.

Finally, we consider risks and failure modes observed. For example, what if the map wrongly predicts a lower speed and the driver slows unnecessarily? In our tests, a false slow down might annoy the driver or, worse, pose a rear end risk if traffic flow is faster. We will check that no dangerous false alarms happen if our system had any, that’s a red flag (this likely only if map errors were severe; our conflict resolution should catch obvious ones). Another aspect: if connectivity is lost (no live map access), does the system gracefully degrade? We test a scenario driving offline the system should then rely on cached map data or just behave like normal TSR. Ensuring that even in failure, the system doesn’t mislead the driver is key (for instance, if map can’t be verified, maybe don’t give early warning rather than give a wrong warning).

Through these evaluations, we aim to demonstrate the advantage of a map augmented approach and provide guidance on which map sources, or combinations of sources, yield the best performance. We expect the results to show that a fused strategy—using multiple maps together ≥with continuous vali- dation—can meet the stringent accuracy requirements (99% correct) and timeliness targets (warnings issued in > 90% of cases at least 3 s earlier than camera only detection) mandated by emerging ISA standards. [11]. The experiments will also highlight areas for improvement, such as the importance of keeping maps updated (perhaps via crowd inputs) and the need for robust localization in certain environments..

7. Conclusion: Feasibility, Limitations, and Outlook

The research presented shows that leveraging smartphone accessible map data to predict speed limit changes is not only feasible but highly beneficial for overcoming the inherent limitations of camera only systems. By fusing digital maps with improved vehicle localization (using BLE/ Wi-Fi) and on board vision, an intelligent speed assistance system can anticipate speed zone transitions with greater lead time and reliability. This can enhance driver comfort (smooth transitions instead of hard braking) and safety (ensuring the vehicle is at the correct speed before a legal limit change, even if a sign is missed or obscured).

Real world Feasibility: Many new vehicles already have the necessary pieces navigation maps with speed data, forward cameras, and connectivity. Our architecture mainly requires integrating these and adding perhaps a phone based augmenta- tion. For older vehicles or aftermarket, a smartphone alone can serve as the platform (using its GPS, camera, and connectivity) though interfacing with vehicle speed would improve it (via OBD II Bluetooth dongle). Cloud services from Google, HERE, etc., make the data accessible, and crowdsourced projects (World Bank, etc.) are rapidly expanding open speed limit databases in countries lacking them [9]. The trend in regulation (e.g. ISA now mandatory in new EU cars [11]) means vehicles will either have to rely on cameras + maps or a combination to meet accuracy requirements. Our approach directly aligns with these needs by using multi source data fusion to hit the accuracy targets (we saw that camera or map alone might not hit 90% reliability, but together they can exceed it [11]).

Limitations and Risks: Despite the promise, there are chal- lenges to address. Data reliability and timeliness: If map data is outdated or incorrect, the system could either warn unnecessarily or, worse, fail to warn when needed. We mitigate this via the feedback loop, but that relies on enough vehicles contributing data. In sparse usage, a wrong map entry might persist. There’s also a risk of over reliance on automation drivers might become complacent if the system usually warns them. If one day both camera and map fail to detect a new limit (rare but possible), the driver could be caught off guard. Thus, human factors design (the system should encourage driver to still pay attention to signs) is important. Connectivity dependence: Using live APIs (Google, etc.) means the system needs internet. In tunnels or remote areas without cell signal, the system should fall back gracefully (e.g. use last known data or just camera). Caching maps locally can help, but smart- phone apps have limits on storage and freshness. Localization privacy: Scanning Wi-Fi/BLE and uploading observations to cloud might raise privacy questions (though it’s similar to what phones already do for location services). Proper anonymization and security must be in place if sharing data. Infrastructure: Our proposal works best if augmented by some infrastructure (beacons in tunnels, up to date city maps). Not all regions will invest in that, so the system in those places might revert to lower performance (still better than camera alone, but not as good as in a fully connected city).

Safety considerations: If the system is tied into vehicle control (active ISA), any errors could directly affect vehicle behavior. A map error causing the car to slow unexpectedly on a highway could pose a hazard. We addressed this by conservative design (prefer false slow down over false speed up, and include driver override at all times). Extensive testing and perhaps geofencing of where active control is allowed might be needed (e.g. only actively intervene when very confident in data). Cybersecurity also comes up: a malicious actor shouldn’t be able to spoof map data or sensor inputs to trick cars into inappropriate actions secure channels and validation for map updates are necessary (which industry is indeed working on for OTA updates with signatures).

Outlook: The integration of maps and sensors for ISA is a stepping stone toward more advanced autonomous driving maps. In the future, real time crowdsourced mapping (so called “live maps”) will likely eliminate much of the stale data problem [20]. Vehicles might share sign detections with the cloud in real time (some upscale cars do this already for traffic signs, sending info back to manufacturers). This will ensure map speed limits are always up to date, effectively merging the map and camera into one cooperative system. Furthermore, with V2X communications, infrastructure might broadcast speed limits and changes ahead (some corridors in EU are testing sending speed limit info directly to cars). Our system could easily ingest such broadcasts as another source,further improving reliability.

Another development is using predictive analytics on speed limits: e.g. if many cars begin slowing at a certain location regularly, an AI could predict there's likely a speed reduction (even before a formal map update or sign detection, almost like Waze reports but automated). This could help with things like dynamic school zone detection (cars slowing at 8 am near a school implies the 20 mph zone is active, so warn others even if they haven't reached the sign).

Conclusion: In conclusion, augmenting ADAS with smart- phone derived map data and enhanced localization is a power- ful approach to achieve early and reliable speed limit aware- ness. Our literature review and experimental findings support that camera based recognition, while useful, is not sufficient alone under all conditions [5], [15]. Map data provides the necessary context and foresight [11], and when combined with real time sensor validation [6], the result is an ISA system robust against occlusion, weather, and signage issues. This fu- sion approach meets the emerging requirements for intelligent speed assistants and can be practically deployed using today's connected technologies. By addressing data accuracy through continuous feedback, the system also contributes to the broader goal of self-healing maps that stay current with the roads. Ultimately, this synergy between digital maps and vehicle perception moves us closer to safer roadways where the posted speed limit is always known and respected, protecting both drivers and vulnerable road users.

Acknowledgments

The authors acknowledge the contributions of industry practitioners, regulatory bodies, and open access researchers whose datasets and evaluation frameworks have informed this work. The authors acknowledge the use of OpenAI (GPT-5 model, 2025) and Grammarly to assist in improving the clarity, organization, and structure of this manuscript.

Conflict of Interest

The author declares that there are no conflicts of interest related to the publication of this paper. The author conducted this research entirely independently of their professional duties and responsibilities at their respective employing organization. The research and opinions presented in this paper are solely those of the authors and do not represent the views, positions, or policies of their respective employers or affiliations.

References

[1] S. Miyata, "Automatic recognition of speed limits on speed-limit signs by using machine learning," *Journal of Imaging*, vol. 3, no. 3, p. 25, 2017.

[2] A. F. Magnussen, N. Le, L. Hu, and W. E. Wong, "A survey of the inadequacies in traffic sign recognition systems for autonomous vehicles," *International Journal of Performability Engineering*, vol. 16, no. 10, p. 1588, 2020.

[3] S. Kumar, T. Brophy, E. M. Grua, G. Sistu, V. Donzella, and C. Eising, "Evaluating the impact of weather-induced sensor occlusion on bevfu- sion for 3d object detection," *arXiv preprint arXiv:2511.04347*, 2025.

[4] J. Xing, M. Nguyen, and W. Qi Yan, "The improved framework for traffic sign recognition using guided image filtering," *SN Computer Science*, vol. 3, no. 6, p. 461, 2022.

[5] M. Seraj, A. Rosales-Castellanos, A. Shalkamy, K. El-Basyouny, and T. Z. Qiu, "The implications of weather and reflectivity variations on automatic traffic sign recognition performance," *Journal of advanced transportation*, vol. 2021, no. 1, p. 5513552, 2021.

[6] H. Jamshidi, T. Lukaszewicz, A. Kashi, A. Berghuvud, H.-J. Zepernick, and S. Khatibi, "Fusion of digital map traffic signs and camera-detected signs," in *2011 5th International Conference on Signal Processing and Communication Systems (ICSPCS)*. IEEE, 2011, pp. 1–7.

[7] H. B. Fredj, A. Chabbah, J. Baili, H. Faiedh, and C. Souani, "An efficient implementation of traffic signs recognition system using cnn," *Microprocessors and Microsystems*, vol. 98, p. 104791, 2023.

[8] "Speed limits — roads api — google for developers," [Online; accessed 2025-11-29]. [Online]. Available: https://developers.google.com/maps/documentation/roads/speed-limits

[9] "How can we help cities collect better speed limit data and protect road users?" [Online; accessed 2025-11-29]. [Online]. Available: https://blogs.worldbank.org/en/transport/intelligent-speed-assistance0

[10] S. Shashidharan, "How ai and imagery keep speed limits on google maps updated," 10 2023, [Online; accessed 2025-11-29]. [Online]. Available: https://blog.google/products/maps/ how-ai-and-imagery-keep-speed-limits-on-google-maps-updated/

[11] TomTom, "Adas map speed limits," [Online; accessed 2025-11- 29]. [Online]. Available: https://download.tomtom.com/open/banners/Whitepaper-ADAS-Map-Speed-Limits.pdf

[12] "ios and ipados 26 - feature availability - apple," [Online; accessed 2025-11-29]. [Online]. Available: https://www.apple.com/ios/ feature-availability/#maps-speed-limits#

[13] R. Floman, "Maps vs. ground truth in phoenix, az: Ready for autonomous driving?" 3 2022, [Online; accessed 2025-11-29]. [Online]. Available: https://data.getnexar.com/blog/phoenixspeedlimit/

[14] C.-W. Ang, "Vehicle positioning using wifi fingerprinting in urban environment," in *2018 IEEE 4th World Forum on Internet of Things (WF-IoT)*. IEEE, 2018, pp. 652–657.

[15] H. Kandula, V. Chidurala, Y. Cao, and X. Li, "Localization of a ble device based on single-device rssi and doa measurements," *Network*, vol. 4, no. 2, pp.

196–216, 2024.

[16] P. J. Garc´ıa-Paterna, A. S. Mart´ınez-Sala, and J. C. Sa´nchez-Aarnoutse, “Empirical study of a room-level localization system based on bluetooth low energy beacons,” *Sensors*, vol. 21, no. 11, p. 3665, 2021.

[17] Y. Su, L. Chen, and X. Liu, “Crowdsourced wifi fingerprint localization in urban canyon,” *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 46, pp. 177–184, 2022.

[18] M. Ibrahim, A. Rostami, B. Yu, H. Liu, M. Jawahar, V. Nguyen,M. Gruteser, F. Bai, and R. Howard, “Wi-go: accurate and scalable vehicle positioning using wifi fine timing measurement,” in *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services*, 2020, pp. 312–324.

[19] K. Nur, S. Feng, C. Ling, and W. Ochieng, “Integration of gps with a wifi high accuracy ranging functionality,” *Geo-spatial Information Science*, vol. 16, no. 3, pp. 155–168, 2013.

[20] R. Floman, “Maps vs. ground truth in phoenix, az: Ready for autonomous driving?” 3 2022, [Online; accessed 2025-11-30]. [Online]. Available: https://data.getnexar.com/blog/phoenixspeedlimit/

[21] F. Jime´nez, J. E. Naranjo, J. J. Anaya, F. Garc´ıa, A. Ponz, and J. M. Armingol, “Advanced driver assistance system for road environments to improve safety and efficiency,” *Transportation research procedia*, vol. 14, pp. 2245–2254, 2016.