

Multiagent Sensor Fusion for Connected & Autonomous Vehicles to Enhance Navigation Safety

Suryansh Saxena, Isaac K. Isukapati, Stephen F. Smith, John M. Dolan
(suryanss, isaack, sfs, jmd)@cs.cmu.edu

The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA - 15213

Abstract. Today, autonomous vehicle (AV) navigation systems rely solely on local sensor data feed for safe & reliable navigation. However, it is not uncommon for sensor data to contain erroneous measurements resulting in false predictions, classified as either false positives (predict non-existent obstacle) or false negatives (e.g., missed obstacle). In this paper, we propose a methodology to identify and minimize false negatives in autonomous vehicle navigation, since these are arguably the most dangerous. According to the methodology, each autonomous agent simultaneously localizes and maps its local environment. This map, in turn, is encoded into a low-resolution message and shared with nearby agents via DSRC, a wireless vehicle communication protocol. Next, the agents distributively fuse this information together to construct a world interpretation. Each agent then statistically analyzes its own interpretation with respect to the world interpretation for the common regions of interest. The proposed statistical algorithm outputs a measure of similarity between local and world interpretations and identifies false negatives (if any) for the local agent. This measure, in turn, can be used to inform the agents to update their kinematic behavior in order to account for any errors in local interpretation. The efficacy of this methodology in resolving false negatives is shown in simulation.

Keywords: False Negatives, Autonomous Driving, Maximum Deviation Test, Connected Vehicles, DSRC, Sensor Sharing, Sensor Fusion.

1 Introduction

Safety and reliability are the paramount goals of autonomous vehicle (AV) navigation systems, but contemporary AV systems face critical obstacles along the road to attaining these goals. One such obstacle is ubiquitous reliance on a data feed from local sensors, and the accuracy of those data. It is not uncommon for sensor data to contain erroneous measurements that reduce overall safety and confidence in navigation. Erroneous measurements may be a result of failing sensor health, sensor drift, bad calibration, and/or temporary conditions such as inclement weather. Those data may subsequently produce false predictions, which can be classified as either false positives or false negatives. In the context of AV navigation, a false positive arises when vehicle sensors predict a non-existent obstacle, whereas a false negative may manifest in the form of a missed obstacle, an imperfect reading of the road lane, an incorrect speed estimation of other vehicles, etc. Though false positives may influence safety in an indirect manner,

false negatives significantly affect navigational safety. This observation is supported by the fact that the majority of autonomous car accidents reported today are due to false negatives [1] [2].

One example of the direct consequences of this problem occurred in February 2017. In that instance, erroneous sensor readings caused Uber’s self-driving vehicle to miss several traffic signals and stop signs during a test run in San Francisco [3]. While it is true that there are significant efforts in the scientific community to use vision-based machine learning algorithms that will identify and isolate false negatives by recognizing and classifying environmental scenes [4] [5] [6], these methodologies depend heavily on extensive and exhaustive training for reasonable performance. More importantly, these methods do not directly address issues related to sensor drift or environmental noise.

Sensor drift is another major contributing factor to false negatives in local interpretations. Traditionally, issues related to sensor drift have been addressed by online-sensor calibration routines, which try to correct for persistent errors [7] [8] [9] on an ex-post-facto basis. These techniques usually depend upon more robust sensors, such as odometers, to calibrate the others. However, there are two major shortcomings with these methodologies: First, since calibration occurs in a dynamically changing environment, the reliability of correction degrades drastically in certain cases, and the routine must be repeated. Second, these routines have little or no ability to correct errors that manifest due to poor environmental conditions. Since it is difficult to quantify the source of error in dynamic navigation, some of the techniques presented are incomplete solutions to the problem of false negatives, and false negatives remain difficult to detect. For AV driving and navigation to succeed, it is critical to identify and minimize the number of false negatives.

In this paper, we explore a methodology for identifying and minimizing false negatives in local environment interpretation by sharing and fusing sensor data collected by close-proximity autonomous or intelligent vehicles. For each agent, the methodology compares the local obstacle maps with maps generated by other close-proximity agents to identify false negatives in local interpretation.

According to the methodology, each autonomous agent simultaneously localizes and maps its local environment. This map, in turn, is encoded into a low-resolution message and shared via Dedicated Short Range Communication (DSRC), a wireless vehicle communication protocol. Next, the agents distributively fuse this information together to construct a world interpretation. Each agent then statistically analyzes its own interpretation with respect to the world interpretation for the common regions of interest. The proposed statistical algorithm outputs a measure of similarity between local and world interpretations and identifies false negatives (if any) for the local agent. This measure, in turn, can be used to inform the agents to change their kinematic behavior in order to account for any errors in local interpretation. Finally, each agent records the measure and instances of erroneous interpretations, which improves the analysis and quantification of sensor health over time.

As mentioned above, the information is shared between and among vehicles via DSRC signals. DSRC is a two-way short-to-medium-range wireless communications capability that permits very high-frequency data transmission critical to active com-

munications based safety applications. DSRC operates on a dedicated frequency of 5.9 GHz that guarantees communication latencies in the 100-millisecond range. The DSRC protocol provides infrastructure for connected vehicles and enables sensor data sharing between the vehicles. The SAE J2735 [10] standard provides guidelines for exchange of safety-critical data between vehicles (V2V), and between vehicles and infrastructure (V2I) [11]. We chose DSRC for this work because it provides a robust, low-latency dedicated communication infrastructure. However, the proposed methodology operates independently of the DSRC protocol. Any low-latency, high-reliability communication protocol can be used with the proposed methodology to achieve identical performance.

While on the surface, our methodology bears a resemblance to distributed SLAM techniques, a careful analysis highlights the differences. For example, authors in [12] [13] [14] [15] propose the notion of using multi-agent and distributed computing techniques for map-merging. The main objective of these research efforts is to build local maps in a computationally efficient manner. Similarly, authors in [16] [17] [18] [19] explore the ideas of multi-vehicle SLAM techniques for scenarios under which a single sensing platform may not be sufficient for collecting data or creating maps of an unknown environment. Furthermore, the authors in [20] [21] propose nonlinear optimization techniques to minimize errors in local map data association. In specific, they treat errors in local maps in data association as an assignment problem rather than a sensor measurement error problem. Lastly, authors in [22] propose a SLAM technique specifically designed to address the communication and computational issues that affect multi-robot systems. None of these methods takes into consideration the uncertainties associated with local sensor errors.

Thus the framework presented in this paper bears the following distinctions:

1. It takes into consideration the measurement errors in sensing (either due to sensor drift or uncertainties in the environment)
2. It proposes a technique - called the Maximum Deviation Test (MDT) that is capable of highlighting statistical similarities (or differences) between two PDFs in a computationally efficient manner

The rest of the paper is organized as follows: Section 2 presents details of the methodology, Section 3 describes the details of simulation experiments, Section 4 discusses the performance characteristics of the algorithm, and Section 5 provide conclusions and point out future lines of work.

2 Methodology

As mentioned in Section 1, we explore and develop a methodology for sharing and fusing sensor data between multiple autonomous vehicles that have overlapping views of the environment. The methodology focuses on identifying and minimizing the number of false negatives in sensing and interpretation.

Before presenting details of the methodology, it is important to discuss modeling assumptions made to ease some constraints: First, the road network and environment conditions are such that false positives do not have a significant impact on safety. Thus in this work we only identify and minimize false negatives. Second, the vehicles must

have compatible sensors or algorithms, such that the data from multiple agents can be fused without considering format or synchronization inconsistency (as this is not the focus of the paper). Lastly, we assume that only a minority of autonomous vehicles have strong noise associated with their data. In other words, this paper does not address pathological cases like when data from all the agents are highly erroneous.

The methodology can be divided into the following steps: 1) Localization, Segmentation, and Super Frames; 2) Data Fusion along with Local vs Global interpretation. The subsequent text in this section provides further details on each step.

2.1 Localization, Segmentation, and Super Frames

An autonomous agent must first understand its surrounding and its current location for the purpose of navigation. Today, autonomous vehicles are equipped with local sensor systems like LiDAR and camera to build an interpretation of the surrounding world. This is done by scanning different obstacles and landmarks within the field of view. Usually, a LiDAR returns a 360-degree scan of the obstacle field which in turn can be converted into a Cartesian map of the points detected by the LiDAR. In turn, data segmentation algorithms are used to associate a group of points with a particular obstacle. In this paper, we simulate a LiDAR sensor to generate a obstacle grid and we cluster the data based on the sequential compatible nearest neighbor (SCNN) approach [23].

The resulting map after the data segmentation step is translated into an occupancy grid that encodes the estimated distances to the obstacles, and the measure of confidence associated with them. We utilize the Extended Kalman Filter [24] algorithm to track and estimate the obstacle locations. The output of a EKF tracker is the pose estimation for the different segmented obstacles which can be transformed into a map that encodes the distance to all perceivable obstacles and a measure of the accuracy in form of variance into an occupancy grid.

Every EKF cycle, the agents record the estimated location of each obstacle, and the associated variance into a frame. The information in this frame along with the vehicle's speed and GPS location are in turn formatted into a DSRC SAE J2735 message, and shared over DSRC communication channel. We refer to this low-resolution message as a *super frame*. In this paper, the agents share super frames at a frequency of 1 Hz using the SAE J2735 Basic Safety Message(BSM) Part 2 structure. The BSM-2 encodes local vehicle kinematics along with a low resolution super-frame.

2.2 Data Fusion

Autonomous vehicles share and receive super frames every second. Please note that vehicles can only receive super frames from other autonomous vehicles that are in the DSRC range (1000 feet). As mentioned earlier, the rationale behind sharing super frames is that the agents can compare and validate local interpretations of their surroundings with those of others. An individual agent must first identify regions of interest in its field of view that overlap with that of other agents as a comparison can be made only in regions that are observable to the other agents.

Furthermore, to account for different orientations of agents the data from super frames should be transformed into a common frame of reference. Therefore, the data

received from different super frames are transformed into the coordinate frame of the local agent. Euclidean translation and rotation transformation is used for this purpose.

Once the data points have been transformed, we must fuse each data point from the different interpretations from close-proximity agents to the data points from the local vehicle. This data association is usually done using the K-nearest neighbors algorithm, where we cluster the data points based on mean distance and relative angles. The output lets the agent map its own local data points to points in the interpretations from the different agents. The data points in each cluster, except the data points from the local view, are fused with each other using a kernel mixture model. Each data point is basically a probability density function (PDF) whose mean is the distance from the local sensor.

To correct for false negatives, each agent then statistically analyzes its own interpretation with respect to the world interpretation (fused data) for the common regions of interest. A Maximum Deviation Test (MDT) is used for this purpose [25].

Algorithm 1: Maximum Deviation Test

```

 $(\mu_l, \sigma_l)$  = mean and sd of local interpretation;
 $(\mu_g, \sigma_g)$  = mean and sd of world interpretation;
 $\delta_{tol}$  = error tolerance threshold;
Let  $F(x)_l, F(x)_g$  be CDFs for local and global interpretations;
s = 0 (initiate test score);
for p in [0,100] do
     $\delta = \frac{F_l^{-1}(p) - F_g^{-1}(p)}{F_l^{-1}(p)} * 100$ ;
    if  $abs(\delta) \leq \delta_{tol}$  then
        | s += 1
    end
end
if  $s \geq s_{min}$  then
    | return density functions are statistically similar
end

```

As the name suggests, the maximum deviation test is a statistical technique to quantify statistical differences between two density functions. The methodology employed here first generates cumulative density functions (CDFs) for local and global interpretations, and then generates a test score that measures statistical similarity. Here the test-score is nothing but the number of percentile values in a local interpretation CDF that are within user-defined tolerance bounds from the global interpretation. If the test score is less than a preset threshold, then it can be inferred that local and global interpretations are statistically significantly different, suggesting the presence of a false negative. Hence through the MDT we are able to track false negatives in autonomous navigation. Pseudo-code for the methodology is given in Algorithm 1:

An example where the test indicates a false negative is presented in Figure 1. The blue curve corresponds to the data interpretation from the local agent. In this example, the δ_{tol} is set to 5%, and s_{min} is set to 95. The MDT test score turns out to be 85, indicating the presence of a false negative.

Most non-parametric tests, such as the Kalmagorov Smirnov (KS) test, use maximum deviation from the mean as a measure to check for dissimilarity. Therefore, these




Fig. 1. MDT identifies a false negative

tests fail to recognize dissimilarities in heavy-tailed, or multi-modal distributions. On the other hand, MDT uses sum of deviations of every percentile on the distribution as a measure for dissimilarity. This property in addition to the symmetric nature of the test makes MDT a very powerful test over either the KS Test or the Kullback-Leibler (KL) Divergence test. Therefore, it is appropriate to use MDT for comparing occupancy grid maps data between close-proximity autonomous agents.

3 Simulation Framework

Testing the efficacy of the proposed methodology in simulation raises three broad requirements. First, the simulator should be able to implement microscopic traffic flow characteristics of individual vehicles (e.g., position, velocity, car-following, and lane changing behavior). Second, the simulator should be able to simulate the DSRC communication protocol. Third, simulator should be able to simulate the behavior of autonomous vehicles, which includes generating LiDAR scans, LiDAR data segmentation with respect to agents, building obstacles maps and finally comparing the fused view with respect to the local view via the Max Deviation Test. No existing simulator satisfies all three requirements. However, the combined capabilities of the open source simulators SUMO and ROS do satisfy all three requirements. Therefore, we developed a software package that provides an interface between SUMO's microscopic traffic simulator and ROS.

SUMO can generate a traffic networks, implement traffic rules, and manage and maintain microscopic traffic flow characteristics. Furthermore, the Veins library in SUMO simulates the DSRC communication protocol. Moreover, the behavior of traffic objects inside SUMO can be accessed and manipulated through the TraCI API. This feature is very critical for simulating and controlling autonomous vehicle behavior. ROS is an ideal choice because it has useful repositories for simulating autonomous behaviors that are reflective of the real-world.

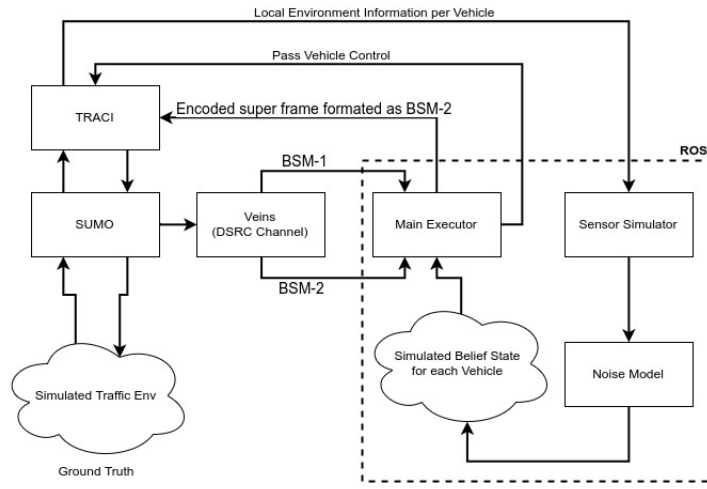


Fig. 2. High-Level Schematic of the Simulator

The software architecture is detailed in Figure 4. As can be seen, SUMO sets up the traffic and vehicular environment and updates vehicle motion models at each simulation step. Some of the vehicles in SUMO are treated as autonomous. TraCI generates an environment (map) grid to simulate local sensor data. TraCI simulates a 360-degree Velodyne LiDAR that generates the occupancy grid for each autonomous vehicle. It should be noted that TraCI passes the local state and sensor information for each autonomous vehicle through a noise model to emulate real-world sensor data that are generated by an autonomous vehicle. The Map grid data generated are used by each agent to develop local obstacle maps, which are later processed into super frames. The TraCI API, also sends information related to other autonomous vehicle's locations and instantaneous kinematics are passed to ROS via local DSRC channel. Lastly, super frames are created for each autonomous vehicle and shared with other autonomous vehicles within DSRC range, through the DSRC channels via TraCI.

3.1 Simulator Assumptions

While it is our main objective to minimize the disparity between the simulator and the real world, the simulation framework makes a few assumptions to relax certain constraints. The assumptions made are as follows:

- All interactions within the simulation are event-based. That is, at a certain time-step an event is initiated and agents appropriately interact with the simulator. This assumption was made to relax the issue of clock synchronization, which is beyond the scope of this work.
- No false positives exist in sensor readings.
- All vehicles in the simulator can travel only at speeds less than or equal to 80 mph.
- Only other vehicles and road geometry are qualified as obstacles.

- All noise in the simulation is uni-modal. We intend to extend the framework to multi-modal systems in future.
- There is low noise in GPS, Odometer and IMU data. This assumption can be easily removed by implementing efficient multi-modal localization software; however, this was not the interest of this research.
- There are no elevations or depressions in the road network and the height of each vehicle is the same. These assumptions resolve into modulating the LiDAR sensor model with one LIDAR beam for building 2D obstacle maps. It is not difficult to deal with multi-beam LiDAR and build a more accurate and complex occupancy grid, but that is not the interest of this research.
- All autonomous vehicles are connected. We also assume in this work that corresponding network dead-regions are quite sparse and don't significantly affect the communication.

Algorithm 2: Simulator Process Flow

```

 $\delta_t = 100$  millisecond (initialize time step);
 $t = 0$  (initiate simulation time);
 $A_t =$  set of all AVs in the system during time  $t$ ;
 $a_i =$  autonomous vehicle 'i';
 $A_j =$  set of AVs within DSRC range of  $a_i$ ;
 $T =$  time when all the vehicles exit the simulation;
while  $t \leq T$  do
    freeze simulation frame ;
    for  $a_i \in A_t$  do
        TraCI sends local sensor feed to ROS;
        builds obstacle maps using EFK;
        update kinematic model;
        encode location and speed info into BSM-1 ;
        ROS sends back BSM-1 to SUMO via TraCI;
        receives BSM-1 via Viens from all  $a_j \in A_j$ ;
        if  $t \% 1000 == 0$  then
            ROS encodes occupancy grid into super-frame;
            ROS sends super frame to SUMO via TraCI;
            receives BSM-2 via Viens from all  $a_j \in A_j$ ;
            process message(s) to develop world interpretation;
            compares local interpretation with fused interpretation;
            identify false negatives and update kinematic model;
        end
    end
    SUMO updates collision model & reports potential collision  $\forall a_i \in A_t$ ;
    unfreeze simulation frame;
    sleep( $\delta_t$ );
     $t += \delta_t$ ;
end

```

3.2 Experimental Design

The objective of the proposed methodology is to reduce false negatives in sensing and to enhance the safety and reliability of autonomous vehicle navigation. The methodology's

efficacy is best tested in a high-risk accident-prone environment. It is generally accepted that the combination of high-speed merges along with blind spots makes lane changing on a highway highly accident-prone [26].

We designed an experiment to simulate lane changing behavior on a 2-mile-long straight highway with 3 lanes. Furthermore, the freeway segment has three on and off ramps located equidistant from one other. The following three scenarios are considered:

- Scenario - 1: Total cars = 50; % AVs = 50
- Scenario - 2: Total cars = 100; % AVs = 50
- Scenario - 3: Total cars = 200; % AVs = 50

Three cases were created for each scenario with the percent of AVs with erroneous sensors set to 5%, 10%, and 20%.

For a given scenario, and a case, it can be inferred that the number of cars in the system, % AVs, and % AVs with erroneous sensors is constant. One needs to keep two objectives in mind to ensure thorough testing: 1) vary vehicle arrival pattern for a given input volume; and 2) for a given vehicle arrival pattern distribute the location of autonomous cars, and AVs with erroneous sensors. To meet the first objective, we ran five Monte Carlo simulations with five random seeds and to address the second objective, ten Monte Carlo simulations are run to uniformly distribute location of AVs, and AVs with erroneous sensors. Hence, results for a given scenario, and cases are aggregated over 50 Monte Carlo simulations. Algorithm 2 presents pseudo-code for the simulation process

4 Analysis of Results

As mentioned earlier, this research proposes a methodology that identifies and minimizes false negatives in autonomous vehicle navigation by sharing and fusing sensor data of close-proximity autonomous or intelligent vehicles. Simulation experiments are designed to evaluate the effectiveness of this methodology and this section focuses on the analysis of the results. However, before proceeding any further, we will first define what successful resolution of a false negative means in the simulation. As stated before, ROS controls vehicle kinematics for autonomous vehicles while SUMO controls vehicle kinematics of non-autonomous vehicles [27].

Every simulation step, path planning for an individual autonomous vehicle is done in ROS using its local environment interpretation, and after correcting for any false negatives. These path planning decisions are passed to SUMO via TraCI for implementation. In turn, SUMO cross-validates these control decisions and makes corrections in case of an impending collision. In that sense, if SUMO implements vehicle path planning decisions without making any adjustments, then there are no inconsistencies in the autonomous vehicle's perception of its surroundings; otherwise, it can be inferred that the algorithm failed to identify and correct false negatives. Therefore, every simulation step, the following information is logged for every autonomous vehicle in the system: 1) number of vehicles it interacted with for building the collaborative world interpretation; 2) a binary indicator value for any false negative resolution; 3) path planning decisions computed using local sensor data and 4) SUMO's path planning decisions.

We post-processed simulation log files to compute the number of instances in which the proposed algorithm was able to successfully resolve false negatives. The results are summarized in Table 1, Table 2 and Table 3 and presented in Fig 3. Each table has fourteen columns. Values in columns 1-2 represent the number of vehicles in the simulation, and proportion of AVs with faulty sensors. Values in columns 3-14 represent the number of instances per simulation a false negative was corrected, % of times the proposed methodology successfully corrected a false negative, and % of times it failed to correct a false negative. Furthermore, results are subdivided into number of the autonomous vehicles involved in building the world interpretation (2, 3, 4, and 5 or more vehicle interactions). Table 1 presents results for 5% autonomous vehicles with faulty sensors. Similarly, Table 2 and Table 3 present results for 10%, and 20% AVs with faulty sensors. Based on the values presented in these tables, it is easy to see that the proposed algorithm successfully corrected false negatives about 95-99% of times the in case of 5 or more autonomous vehicle interactions. These values are between 89-96% for four vehicle interactions, 82-92% for three vehicle interactions, and 72-82% in case of two vehicle interactions.

Table 1: Observations for scenarios with 5% of autonomous vehicles with faulty sensors

Total # of vehicles	% Avs with faulty sensors	Avg. 2-vehicle interactions			Avg. 3-vehicle interactions			Avg. 4-vehicle interactions			Avg. 5 or more vehicle interactions		
		Total	% Success	% Fail	Total	% Success	% Fail	Total	% Success	% Fail	Total	% Success	% Fail
50	5	26	77.2	22.8	42	86.2	13.8	22	93.4	6.6	21	98.5	1.5
100	5	36	79.4	20.6	46	90.0	10.0	43	95.2	4.8	29	99.0	1.0
200	5	43	82.2	17.8	56	92.2	7.8	64	96.3	3.7	34	99.0	1.0

Table 2: Observations for scenarios with 10% of autonomous vehicles with faulty sensors

Total # of vehicles	% Avs with faulty sensors	Avg. 2-vehicle interactions			Avg. 3-vehicle interactions			Avg. 4-vehicle interactions			Avg. 5 or more vehicle interactions		
		Total	% Success	% Fail	Total	% Success	% Fail	Total	% Success	% Fail	Total	% Success	% Fail
50	10	37	73.5	26.5	34	82.1	17.9	58	91.5	8.5	20	97.0	3.0
100	10	38	74.2	25.8	51	87.5	12.5	74	93.2	6.8	27	98.5	1.5
200	10	38	77.3	22.7	63	89.8	10.2	65	94.0	6.0	49	99.0	1.0

Table 3: Observations for scenarios with 20% of autonomous vehicles with faulty sensors

Total # of vehicles	% Avs with faulty sensors	Avg. 2-vehicle interactions			Avg. 3-vehicle interactions			Avg. 4-vehicle interactions			Avg. 5 or more vehicle interactions		
		Total	% Success	% Fail	Total	% Success	% Fail	Total	% Success	% Fail	Total	% Success	% Fail
50	20	33	70.1	29.9	47	82.0	18	65	89.2	10.8	20	95.3	4.7
100	20	47	72.0	28.0	55	84.8	15.2	85	90.7	9.3	36	96.7	3.3
200	20	63	74.3	25.7	82	85.2	14.8	95	91.0	9.0	50	98.0	2.0

Fig. 3. Summary statistics

The following inferences can be drawn based on these observations: 1) the probability of resolving a false negative increases as the number of autonomous vehicles interactions increase; 2) the variance in resolving a false negative also decreases with increased number of interactions; 3) the probability of resolving a false negative decreases with increased % in autonomous vehicles with bad sensors. These trends can be clearly seen in Figure 4.

5 Conclusions & Future Work

In this paper we presented a methodology that identifies and minimizes false negatives in autonomous vehicle navigation by sharing and fusing sensor data of close-proximity

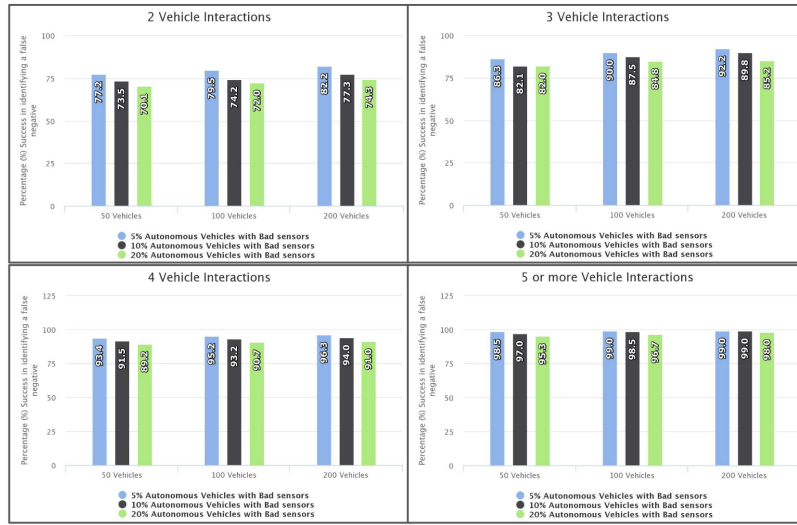


Fig. 4. Percentage success in resolving false negatives for different vehicle group interactions

autonomous or intelligent vehicles. Using the methodology, each autonomous agent simultaneously localizes and maps its local environment. This map, in turn, is encoded into a low-resolution message and shared via DSRC. Next, the agents collaboratively fuse this information together to construct a world interpretation. Each agent then statistically analyzes its own interpretation with respect to the common regions of interest. The proposed statistical algorithm outputs a measure of similarity between local and world interpretations and identifies false negatives (if any) for the local agent. This measure, in turn, can be used to inform the agents to change their kinematic behavior in order to account for any errors in local interpretation.

The efficacy of this methodology is tested in simulation, and based on the simulation results it can be inferred that the methodology is effective in resolving false negatives. We have identified several directions for future work:

- One direction is to extend the framework to identify & minimize false positives.
- Another is to quantify the impact of the percentage of autonomous vehicle penetration on the efficiency of the algorithm (in the current paper, we only considered a penetration level of 50)
- Today, autonomous cars have no obvious way of self-assessing sensor's health. In principle, one can revisit the data logged by the proposed methodology to ascertain sensor's health. This can be done by looking at the number of instances erroneous sensor readings that are corrected.
- In the current paper, equal weight has been given to each obstacle map while generating the world interpretation or fused map. Exploration of fused map creation that takes into account sensor health is another area for future work.

- Finally, it would be interesting to explore multi-agent collaborative path planning. Such systems could have profound impact on improving safety of rural high-speed signalized intersections.

References

1. Rachel Spacek. Tesla car on autopilot warned driver 7 times before fatal crash, safety regulator says.
2. Samuel Gibbs. Google's self-driving car in broadside collision after other car jumps red light.
3. David Z. Morris. Uber's self-driving systems, not human drivers, missed at least six red lights in san francisco.
4. Donald F. Rudny and David W. Sallmann. Analysis of accidents involving alleged road surface defects (i.e., shoulder drop-offs, loose gravel, bumps and potholes). In *SAE Technical Paper*. SAE International, 02 1996.
5. A. Talukder, R. Manduchi, A. Rankin, and L. Matthies. Fast and reliable obstacle detection and segmentation for cross-country navigation. In *Intelligent Vehicle Symposium, 2002. IEEE*, volume 2, pages 610–618 vol.2, June 2002.
6. Henry Schneiderman and Marilyn Nashman. A discriminating feature tracker for vision-based autonomous driving. *IEEE Transactions on Robotics and Automation*, 10(6):769–775, 1994.
7. B.R. Douillard, J.S. Levinson, and G.T. Sibley. Calibration for autonomous vehicle operation, May 4 2017. US Patent App. 14/756,996.
8. Stephan Weiss, Markus W Achtelik, Simon Lynen, Margarita Chli, and Roland Siegwart. Real-time onboard visual-inertial state estimation and self-calibration of mavs in unknown environments. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 957–964. IEEE, 2012.
9. Andrew Howard. Real-time stereo visual odometry for autonomous ground vehicles. In *Intelligent Robots and Systems, 2008. IROS 2008. IEEE/RSJ International Conference on*, pages 3946–3952. IEEE, 2008.
10. SAE International. *Dedicated Short Range Communications (DSRC) Message Set Dictionary*. 2016.
11. Yunxin (Jeff) Li. *An Overview of the DSRC/WAVE Technology*, pages 544–558. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012.
12. Kok Seng Chong and Lindsay Kleeman. Large scale sonarray mapping using multiple connected local maps. In *Field and Service Robotics*, pages 507–514. Springer, 1998.
13. Juan D Tardós, José Neira, Paul M Newman, and John J Leonard. Robust mapping and localization in indoor environments using sonar data. *The International Journal of Robotics Research*, 21(4):311–330, 2002.
14. Michael Bosse, Paul Newman, John Leonard, and Seth Teller. Simultaneous localization and map building in large-scale cyclic environments using the atlas framework. *The International Journal of Robotics Research*, 23(12):1113–1139, 2004.
15. José A Castellanos, Ruben Martinez-Cantin, Juan D Tardós, and José Neira. Robocentric map joining: Improving the consistency of ekf-slam. *Robotics and autonomous systems*, 55(1):21–29, 2007.
16. Stefan B Williams, Gamini Dissanayake, and Hugh Durrant-Whyte. Towards multi-vehicle simultaneous localisation and mapping. In *Robotics and Automation, 2002. Proceedings. ICRA'02. IEEE International Conference on*, volume 3, pages 2743–2748. IEEE, 2002.

17. Diego Rodriguez-Losada, Fernando Matia, and Agustin Jimenez. Local maps fusion for real time multirobot indoor simultaneous localization and mapping. In *Robotics and Automation, 2004. Proceedings. ICRA'04. 2004 IEEE International Conference on*, volume 2, pages 1308–1313. IEEE, 2004.
18. Andrew Howard. Multi-robot mapping using manifold representations. In *Robotics and Automation, 2004. Proceedings. ICRA'04. 2004 IEEE International Conference on*, volume 4, pages 4198–4203. IEEE, 2004.
19. Max Pfingsthorn, Bayu Slamet, and Arnoud Visser. A scalable hybrid multi-robot slam method for highly detailed maps. In *Robot Soccer World Cup*, pages 457–464. Springer, 2007.
20. Alexander Cunningham, Manohar Paluri, and Frank Dellaert. Ddf-sam: Fully distributed slam using constrained factor graphs. In *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*, pages 3025–3030. IEEE, 2010.
21. Alexander Cunningham, Kai M Wurm, Wolfram Burgard, and Frank Dellaert. Fully distributed scalable smoothing and mapping with robust multi-robot data association. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 1093–1100. IEEE, 2012.
22. Maria Teresa Lazaro, Lina María Paz, Pedro Pinies, Jose A Castellanos, and Giorgio Grisetti. Multi-robot slam using condensed measurements. In *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*, pages 1069–1076. IEEE, 2013.
23. Aron Jace Cooper. *A comparison of data association techniques for simultaneous localization and mapping*. PhD thesis, Massachusetts Institute of Technology, 2005.
24. Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic robotics*. MIT press, 2005.
25. Isaac K Isukapati and George F List. Synthesizing route travel time distributions considering spatial dependencies. In *Intelligent Transportation Systems (ITSC), 2016 IEEE 19th International Conference on*, pages 2143–2149. IEEE, 2016.
26. Basav Sen, John D Smith, and Wassim G Najm. Analysis of lane change crashes. Technical report, 2003.
27. Daniel Krajzewicz, Georg Hertkorn, Christian Rössel, and Peter Wagner. Sumo (simulation of urban mobility)-an open-source traffic simulation. In *Proceedings of the 4th middle East Symposium on Simulation and Modelling (MESM20002)*, pages 183–187, 2002.