Localization with Multi-Modal Vision Measurements in Limited GPS Environments Using Gaussian Sum Filters

Jonathan R. Schoenberg, Mark Campbell, and Isaac Miller
Sibley School of Mechanical and Aerospace Engineering
Cornell University, Ithaca, NY 14853
{jrs55, mc288, itm2}@cornell.edu

***Abstract*— A Gaussian Sum Filter (GSF) with component extended Kalman filters (EKF) is proposed as an approach to localize an autonomous vehicle in an urban environment with limited GPS availability. The GSF uses vehicle relative vision-based measurements of known map features coupled with inertial navigation solutions to accomplish localization in the absence of GPS. The vision-based measurements are shown to have multi-modal measurement likelihood functions that are well represented as a weighted sum of Gaussian densities and the GSF is ideally suited to accomplish recursive Bayesian state estimation for this problem. A sequential merging technique is used for Gaussian mixture condensation in the posterior density approximation after fusing multi-modal measurements in the GSF to maintain mixture size over time. The representation of the posterior density with the GSF is compared over a common dataset against a benchmark particle filter solution. The Expectation-Maximization (EM) algorithm is used offline to determine the representational efficiency of the particle filter in terms of an effective number of Gaussian densities. The GSF with vision-based vehicle relative measurements is shown to remain converged using 37 minutes of recorded data from the Cornell University DARPA Urban Challenge (DUC) autonomous vehicle in an urban environment that includes a 32 minute GPS blackout.**

I. INTRODUCTION

Autonomous vehicles provide opportunities to remove humans from operating in dangerous civilian and military scenarios. The autonomy relies on accurate localization of the vehicle in diverse environments, including urban environments where absolute position information is not available from a Global Navigation Satellite System (GNSS). The urban environment is challenging for reliable position estimation from GNSS signals due to multi-path reflections and obstruction of direct path signals. In the absence of absolute position information, autonomous vehicles rely on dead reckoning from an inertial navigation system (INS) to localize the vehicle. Unfortunately, small errors in the INS solution will accumulate into large position deviations after a few minutes, which in turn, prevents the vehicle from localizing itself within a lane.

J. Schoenberg is a Graduate Research Fellow at the Sibley School of Mechanical and Aerospace Engineering, Cornell University.

M. Campbell is an Associate Fellow of the AIAA and Associate Professor at the Sibley School of Mechanical and Aerospace Engineering, Cornell University.

I. Miller is a Postdoctoral associate at the Sibley School of Mechanical and Aerospace Engineering, Cornell University.

The availability of recognizable street markings around an autonomous vehicle motivates the development of map-aided localization techniques. Syed and Cannon [1] present Map-Aided GPS (MAGPS) where road segment information derived from a robust map-matching technique is tightly coupled into the GPS solution as a set of constraints. Also addressing the problem of self-localization, Cui and Ge [2] and Fouque and Bonnifait [3] propose techniques that tightly couple GPS signals with a known map. Cui and Ge strictly restrict motion to the map, while Fouque and Bonnifait treat the road as a noisy measurement. Wijesoma [4] develops localization in the context of simultaneous localization and mapping (SLAM) by constraining the platform to a set of known road segments. All of these algorithms make strong assumptions about correlations between vehicle motion and the known road map. Unfortunately, autonomous vehicles are required to operate free from the strict constraints of road maps and therefore require localization techniques that account for this freedom of motion. The motivating example may be a road block or disabled car that requires the autonomous vehicle to pass on the wrong side of the road, complete an impromptu U-turn, or drive off the road around the obstacle.

The PosteriorPose algorithm presented by Miller and Campbell [5] demonstrates a bootstrap particle filter (PF) to be effective in providing map relative localization in the absence of absolute position measurements. The technique does not constrain the vehicle to the road, but uses vision data of observable features in the known map for localization. The map-aided localization problem is fundamentally multi-modal because the vision data provides measurements of closely-space landmarks with unknown correspondence to the known map. This is similar to the data association problem in SLAM where Bailey and Durrant-Whyte [6] point out that incorrect assignment cannot be reversed and can lead to failure of the SLAM algorithm. The problem solved here differs from SLAM, because we have a known, accurate map that enables prediction of data assignment mistakes that show up as multiple modes in the measurement and posterior densities. Additionally, the state vector is small because of the known map and allows multiple modes to be maintained in the posterior density that would be less feasible in a SLAM approach with a large unknown map.

The PosteriorPose algorithm utilizes the bootstrap PF as

the approximate recursive Bayesian estimation algorithm to handle nonlinear vehicle motion and multi-modal vision-based measurement functions. The PF relies on a point mass representation of the posterior density; Arulampalam *et al.* [7] point out several drawbacks of the PF including finite support over the posterior state space, particle degeneracy, difficulty in selecting an appropriate importance density, and sample impoverishment after resampling. Despite successful real-time implementation in the DARPA Urban Challenge (DUC) [5], the PosteriorPose algorithm suffers from these drawbacks. The Gaussian Sum Filter (GSF) is proposed as a solution to the map-aided localization problem to avoid the drawbacks of the PF while handling the multiple modes of the vision-based measurement and posterior densities.

The Gaussian Sum Filter (GSF) has been used to solve nonlinear recursive Bayesian estimation problems since it was introduced by Sorenson and Alspach [8], who derive the GSF by representing the desired a priori, transition, measurement, and posterior densities as a summation of component Gaussians. The primary problem with the GSF is computation intractability: the number of component Gaussians grows geometrically with each iteration of the GSF. Sorenson and Alspach recognized that a failure to manage the number of terms would limit the utility of the GSF [9], and proposed combining components with equal means and covariances and eliminating terms with neglectable weights. Salmond [10] proposed an iterative mixture component joining technique that is used in this paper to avoid a growing number of terms in the GSF.

Despite the computational challenges of the GSF, Peach [11] and Kronhamn [12] applied the GSF to bearing-only target tracking and each used a pruning approach to manage the number of terms in the posterior density. Similarly, Kwok *et al.* [13] used the GSF to solve the initialization problem in bearing-only SLAM. The mixture reduction was performed by truncation, and component Gaussians were removed based on a sequential probability ratio test (SPRT).

The novelty of this paper is to show Salmond's joining algorithm in a GSF solves the fundamentally multi-modal problem of localizing an autonomous vehicle with a known map in a sparse GPS environment. The algorithm is demonstrated on experimental data recorded from the Cornell University DARPA Urban Challenge (DUC) vehicle [14]. Section II reviews the Gaussian Sum Filter and Salmond's condensation technique. Section III describes the Cornell University DUC vehicle testbed and the specific vision based measurements used for localization. Section IV discusses a technique to analyze the representational efficiency of the particle filter using EM to fit a Gaussian mixture to the particles at each time step. Section V shows the algorithm applied to experimental data in a GPS blackout and compares performance of the PF [5] to the GSF algorithm on the same dataset. Additionally, the posterior density representation for the GSF and PF are compared over the same recorded dataset. Finally, Section VI summarizes with conclusions demonstrating the application of the GSF to ground-based autonomous vehicle localization with multi-modal vision measurements as a better representation of the posterior density than the competing PF algorithm.

II. Localization with Map Relative Measurements

Autonomous vehicles require precise localization for local control and planning. The 2007 DARPA Urban Challenge (DUC) provided a motivating example where autonomous vehicles were required to complete a set of mock supply missions over a 60-mile closed urban course [15]. The vehicle is required to obey all traffic laws, including driving in lanes and following precedent rules at intersections. The autonomous operation demands accurate localization within the map; GPS signal occlusions, reflections, and distortions in the urban environment precludes reliance solely on GPS signals. Cornell's robot approaches these difficulties by combining the Route Network Definition File (RNDF) with local image processing to create vehicle relative measurements that are fused into the localization solution. Cornell's robot takes advantage of stoplines and lane boundaries to determine the vehicle pose $x_k = [e_k; n_k; h_k]$ at time t_k, where e_k and n_k are the vehicle east and north position with respect to the RNDF center, and h_k is the vehicle's heading.

A. Review of Gaussian Sum Filtering

The Bayesian estimation paradigm desires to determine the posterior density of the state x_k given a sequence of measurements $Z^K = \{z_0, \ldots, z_k\}$ from $t = 0$ to $t = t_k$. In the Gaussian Sum Filter (GSF), the posterior probability density $p(x_k|Z^k)$ is approximated as a sum of Gaussian densities [16][9][8]:

$$p(x_k|Z^k) \approx \sum_{i=1}^{N} \omega_k^i \cdot \mathcal{N}\left\{x_k; \hat{x}_k^i, P_k^i\right\} \tag{1}$$

where N is the number of components in the mixture, ω_k^i is the weight associated with the i^{th} Gaussian component, given as the multivariate normal density $\mathcal{N}\left\{x_k; \hat{x}_k^i, P_k^i\right\}$ with mean $\hat{x}_k^i$, and covariance P_k^i. The weights are constrained such that $\sum_{i=1}^{N} \omega^i = 1$. In the limit, as the number of Gaussians N approaches infinity and the covariance P_k^i goes to the zero matrix, the approximation in (1) can be used to represent any probability density with arbitrarily small statistical divergence [8]. To determine the Gaussian components in the posterior density, the extended Kalman filter (EKF) [17] is used.

The vehicle discrete time state space system is modeled via nonlinear process and measurement functions:

$$x_k = f(x_{k-1}, u_{k-1}, v_{k-1}) \tag{2}$$

$$z_k = h(x_k, M) + w_k \tag{3}$$

where in the process model (2), x_{k-1} is the state, u_{k-1} is the input to the system, v_{k-1} is the noise on the measurement of those inputs, called *process noise*, all at time t_{k-1}. In the measurement model (3), z_k is the measurement, x_k is the state and w_k is the measurement noise at time t_k, and M represents the known map (RNDF).

The filter starts at time t_{k-1} with an initial N_g term Gaussian mixture representing the a priori probability density of the state x_{k-1}:

$$p(x_{k-1}|Z^{k-1}) \approx \sum_{i=1}^{N_g} \omega_{k-1}^i \cdot \mathcal{N}\left\{x_{k-1}; \hat{x}_{k-1}^i, P_{k-1}^i\right\} \quad (4)$$

where Z^{k-1} represents any a priori information in the system at initialization. Our development fixes N_g, the number of components in the approximation of the posterior density, as a function of time, but in general this could vary. The prediction and update steps of the GSF [8] lead to a posterior density at t_k with $N_T = N_g N_v N_w$ terms in the Gaussian mixture approximation:

$$p(x_k|Z^k) \approx \sum_{r=1}^{N_g N_v N_w} \omega_k^r \cdot \mathcal{N}\left\{x_k; \hat{x}_k^r, P_k^r\right\} \quad (5)$$

where ω_k^r is the weight on the mixture component, $\{\hat{x}_k^r, P_k^r\}$ is the one-step ahead updated state and error covariance from the EKF [17], N_v is the number of terms in the Gaussian mixture of the transition probability density, and N_w is the number of terms in the measurement likelihood Gaussian mixture .

The MMSE estimate can now be computed for the state and state error covariance as a metric for evaluation, although it is not necessary for the GSF itself:

$$\hat{x}_{k_{\text{MMSE}}} = \sum_{r=1}^{N_T} \omega_k^r \cdot \hat{x}_k^r \quad (6)$$

$$P_{k_{\text{MMSE}}} = \sum_{r=1}^{N_T} \omega_k^r \left[P_k^r + (\hat{x}_k^r - \hat{x}_{k_{\text{MMSE}}})(\hat{x}_k^r - \hat{x}_{k_{\text{MMSE}}})^T\right] \quad (7)$$

The drawback of the GSF is clear: as time moves from t_{k-1} to t_k the number of components in the Gaussian mixture representing the posterior increases from N_g to $N_T = N_g N_w N_v$ components. The number of Gaussian components grows in time for the GSF. Therefore, a condensation technique is required to reduce the number of terms representing the posterior from $N_T = N_g N_v N_w$ to N_g so the cycle can begin again.

B. Gaussian Mixture Condensation

The objective of condensation is to represent a Gaussian mixture with fewer components, but minimize the statistical difference between the full and reduced probability density representations. Salmond [10] proposed an iterative first and second moment preserving merging of components based on joining or clustering. The decision on which components to merge is based on an L^2-norm metric.

The joining technique proposed by Salmond [10] sequentially merges the two components that are most similar as defined by the metric (dropping the time index k for clarity):

$$d_{ij} = \frac{\omega^i \omega^j}{(\omega^i + \omega^j)} \left(\hat{x}^i - \hat{x}^j\right)^T P_{\text{MMSE}}^{-1} \left(\hat{x}^i - \hat{x}^j\right) \quad (8)$$

where P_{MMSE} is the overall mixture covariance computed as in (7). The condensation algorithm proposed in [10] is performed at each time step in the recursive GSF.

III. Autonomous Vehicle Testbed

The experimental data analyzed with the GSF was collected using Cornell University's autonomous Chevrolet Tahoe [14], shown in Figure 1.

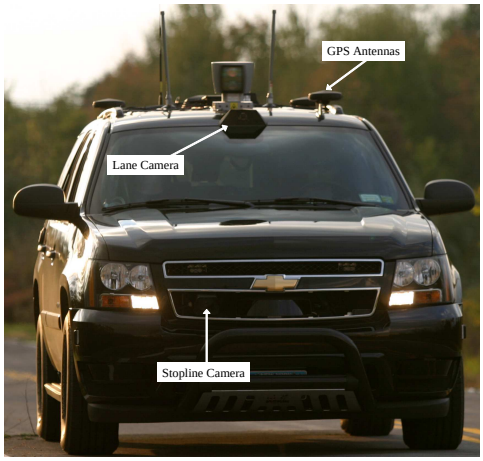

Fig. 1. Cornell University's autonomous Chevrolet Tahoe, equipped with GPS, inertial navigation, and vision-based lane sensing, line sensing, and stopline detection. The Tahoe is shown here on Cornell University's autonomous vehicle test course at the Seneca Army Depot in Romulus, NY.

A. Vehicle Prediction

The vehicle state $x_k = [e_k; n_k; h_k]$ is defined with respect to the map center, where e_k and n_k are the position in the map, and h_k is the heading. The GSF requires predicting the previous posterior density (4) forward to generate the predicted density for measurement fusion. The transition probability density is well represented with a single Gaussian density, $N_v = 1$, because the high update rate of the odometry information results in very small nonlinearities during state prediction.

B. Absolute Position Measurements

The tightly coupled GPS / INS system provides absolute position measurements z_{a_k} that are used to update the vehicle state estimate. The measurement likelihood for the absolute position measurements is well modeled as a single Gaussian given the mean $\bar{x}_k^q$ of the q^{th} component in the predicted Gaussian mixture for $q \in [1, N_g N_v]$:

$$p(z_{a_k}|\bar{x}_k^q) \sim \mathcal{N}\left\{z_{a_k}; \bar{x}_k^q, R_{aa_k}\right\} \quad (9)$$

where R_{aa_k} is the covariance of the absolute position measurement at time t_k. Unfortunately, just as in [5], the absolute position estimates generated via the recursive information filter are correlated from one time step to the next. To account for this autocorrelation, the measurements are whitened by augmenting the state with east and north GPS biases $\beta_k = [\beta_{e_k}, \beta_{n_k}]$ that have the following dynamics:

$$\beta_k = \lambda \beta_{k-1} + \nu_{\beta_{k-1}} \quad (10)$$

where $\lambda = \exp(-\Delta T / T_b)$ accounts for the bias's autocorrelation time T_b during the time interval $\Delta T = t_k - t_{k-1}$, and $\nu_{\beta_{k-1}} \sim \mathcal{N}(0, Q_{\beta\beta})$. The addition of the bias terms now enables the assumption that the likelihood function of the

absolute position information to be Gaussian and white given the q^{th} component of the predicted density:

$$p(z_{a_k}|\bar{x}_k^q,\beta_k^q) \sim \mathcal{N}\left\{z_{a_k};\bar{x}_k^q+\beta_k^q,R_{aa_k}\right\} \tag{11}$$

C. Map Relative Position Measurements

Unlike absolute position measurements, relative position measurements are not straightforward projections of the vehicle position. Instead, the measurements are accurate ranges to nearby stoplines or distances from lane boundaries relative to the vehicle. The relative measurements are combined with the known map to generate weak absolute position information that is fused into the global state estimate. The relative measurements generated from the vehicle cameras allow the vehicle to maintain a global estimate in the absence of absolute position measurements.

1) Stopline Measurements: The first relative position measurement comes from the stopline camera that detects the range from the vehicle to a stopline in the camera field-of-view. The measurement likelihood for the stopline detection is given by a single Gaussian as:

$$p(z_{s_k}|\bar{x}_k^q,M) \sim \mathcal{N}(z_{s_k};\bar{z}_{s_k}^q,R_{ss_k}) \tag{12}$$

where $\bar{z}_{s_k}^q$ is the Euclidean distance from the q^{th} component of the predicted density to the nearest stopline in the map M.

2) Lane Offset Measurements: The second type of relative position measurement comes from the lane finding system that detects bounding lines in the vision image. The measurement z_{o_k} is the perpendicular distance from the detected lane boundaries and the camera heading with respect to the occupied lane. The challenge in incorporating the lane offset measurement is that the vision processing algorithms generate errors that are not well-modeled with a single Gaussian. Instead, the measurement likelihood function is best represented as sum of Gaussians, making it ideal for use in the GSF. The different components of the measurement likelihood come from the different lane detection modes that arise in the vision processing algorithm. There are six different modes of detection by the lane-finding algorithm: detecting the correct lane; detecting the lane to the left of the true lane; detecting the lane to the right of the true lane; the combination of the left and correct lane; the combination of the right and correct lane, and the sixth is the combination of the left, correct, and right lane. The key insight with these uncertainties is that when conditioned on the specific detection mode, the measurement likelihood is well-modeled as a single Gaussian:

$$p(z_{o_k}^l|\bar{x}_k^q,\mu_l,M) \sim \mathcal{N}\{z_{o_k}^l;\bar{z}_k^r,R_{oo_k}^l\} \tag{13}$$

where μ_l indicates the l^{th} measurement mode, and $\bar{z}_k^r$ is the predicted measurement given the r^{th} component mixture for $r \in [1, N_g N_v N_w]$ with respect to the map M. The overall measurement likelihood function for the vision system is given by the Gaussian mixture:

$$p(z_{o_k}|x_k,M) \sim \sum_{l=1}^{N_w=6} \gamma_k^l \cdot \mathcal{N}\{z_{o_k}^l;\bar{z}_k^r,R_{oo_k}^l\} \tag{14}$$

where γ_k^l is the probability of the l^{th} detection mode.

D. Measurement Update and Condensation

The system can now update the N_g term predicted density with absolute position, stopline and lane offset measurements using standard EKF equations [17] and the posterior density (5) is a Gaussian mixture. The algorithm expands the number of components in the Gaussian mixture until condensation is performed. In the implementation discussed here, the high rate odometry data is used as a time synchronizer. The GSF is allowed to run asynchronously between odometry measurements and expands the number of Gaussian components in the posterior density between odometry updates. Before each new odometry update condensation is performed. This allows a variable number of terms to represent the posterior density before condensation.

The number of components in the posterior density after condensation must be less than or equal to the number of components in the original mixture before prediction to avoid an exponentially growing number of Gaussians. The number of components should be enough to accurately represent a multi-modal density, but not too many such that they are representing the same information. For this paper, a fixed number of terms $N_g = 5$ is used to represent the posterior density after condensation. This value was determine by running the filter with $N_g \in [1,10]$ and observing a diminishing performance improvement for $N_g > 5$.

IV. Effective Number of Gaussians

Following condensation, regardless of the technique selected, the GSF approximates the posterior density as a weighted sum of Gaussians. However, the PF implementation represents the posterior density as a discrete set of weighted samples. In an effort to evaluate the accuracy and applicability of the GSF, an off-line fitting procedure is run where the PosteriorPose PF [5] samples at each time step are fit to a Gaussian mixture using the Expectation-Maximization (EM) algorithm [18]. The EM algorithm is initialized using the K-Means clustering technique [18]. The EM algorithm is used along with the Bayesian Information Criterion (BIC) [18] to determine the effective number of Gaussians in the PF posterior samples. This idea of fitting the particle filter data to a Gaussian mixture is similar to Kotecha and Djuric's Gaussian Particle Filter [19] where the particles at each time step are approximated with a Gaussian density that is used for resampling.

V. Experimental Results

A. GSF Localization Performance

The GSF for localization was evaluated using Cornell University's autonomous vehicle driving on a test course. The test course consists of several miles of accurately surveyed roads, accurately painted road lines / stoplines on certain road segments, and segments with no road lines; an overhead view of the course is shown in Figure 2.

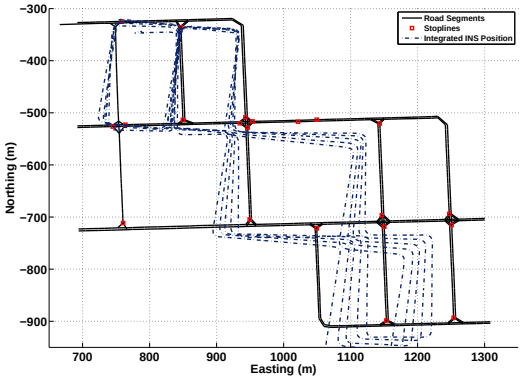

Fig. 2. Overhead view of a portion of Cornell's autonomous test site, which consists of road segments and stoplines. An integrated inertial navigation solution overlaid on the map shows the degradation of vehicle pose estimates in the absence of additional localization information such as GPS or map information.

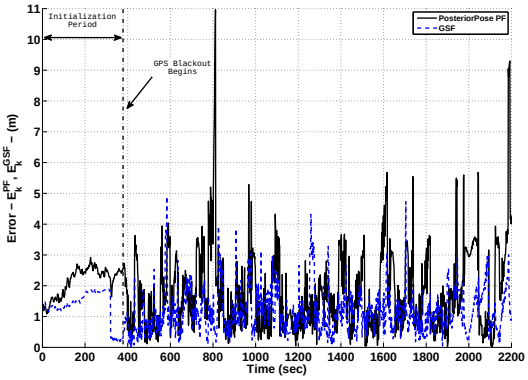

Fig. 3. GSF and Particle filter both remain converged over the entire 37 minute course.

Experimental data was collected for 37 minutes in an area with a clear view of the sky and provided continuous collection of OmniSTAR high precision differential (HP) corrections signal, accurate to 10 cm, which is used as ground truth; neither the PF nor GSF has access to this signal. GPS signals are artificially withheld in order to simulate a GPS blackout. The tests were performed driving at city speeds up to 30 mph (13.4 m/s). The integrated INS solution from this data, shown overlaid on the road network in Figure 2, diverges significantly from the road and demonstrates the need for relative position measurements to maintain accurate global positioning in an extended GPS blackout. To demonstrate the power of relative landmark measurements, the GSF is used to estimate the global east and north positions, these estimates are compared against the truth data to compute position errors, $E_k = ||[e_{HP_k}, n_{HP_k}]^T - [\hat{e}_{GSF_k} - \hat{n}_{GSF_k}]^T||$, at a rate of 10 Hz. All of the analysis included estimating GPS biases to improve the position estimate over the time when the GPS signals are available.

After a 6 minute stationary initialization period during which GPS measurements are available, the GSF estimates the global position with no absolute position measurements. The GSF remains converged during the entire blackout, has a maximum error of 6 m, which occurs after maneuvering through an intersection, and an average error of 1.17 m.

The GSF's state estimates are also compared to Miller and Campbell's [5] PosteriorPose PF algorithm using 2000 particles run on the same data. The position error for each of the algorithms is shown over the entire 37 minute run in Figure 3. There are times where each algorithm performs better than the other, and the PosteriorPose PF algorithm has a maximum error of 11.0 m, again after exiting a turn, and an average error of 1.85 m.

The PF error at each time step E_k^{PF} is substracted from the GSF error at each time step E_k^{GSF} to create a difference in error $\Delta_k = E_k^{\text{PF}} - E_k^{\text{GSF}}$. On average, over the entire run, the GSF performs $\Delta = 0.67$ m better than the PF. When a paired T-test is performed, the 0.67 m average better performance is statistically significant at the $p = 0.05$ level. In fact, the T-test concludes the GSF outperforms the PF over the 37 minute run by at least 0.62 m at the 5% significance level.

An interesting area of comparison between the GSF and PF is when the autonomous vehicle approaches and completes a right turn. This is a particularly challenging case for localization, because of sparse map information and uncertainty associated with the lane of the new road segment the vehicle occupies. Figure 4 shows the area 1970 seconds into the data capture (1590 seconds after GPS blackout).

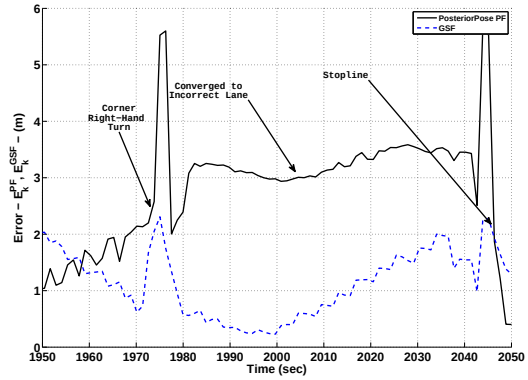

Fig. 4. The autonomous vehicle makes a right turn and the PosteriorPose particle filter converges to the incorrect lane, while the GSF is able to converge to the correct lane. The PF remains diverged until observing a stopline at 2045 seconds.

The PF fails to maintain an accurate representation of the posterior density as the vehicle moves through the corner, and the PF position estimate converges to the incorrect lane after the turn. The PF remains diverged until encountering the stopline at the end of the road segment (62 seconds after exiting the corner).

B. Particle Filter Effective Number of Gaussians

Following the description in Section IV, the EM algorithm is run offline over the stored PF data in order to estimate the effective number of Gaussians in the posterior density at each time step. The EM algorithm is run to convergence for each K-term Gaussian mixture, where $K \in [1, 15]$. The BIC is used

to evaluate which K-term Gaussian mixture best represents the posterior particles. The EM algorithm is run 100 times with different initial means and covariances for each of the K-terms in the Gaussian mixture in an attempt to ensure the algorithm is finding a global minimum at convergence. The number of occurrences where the posterior particles are best represented with a K-term Gaussian mixture is shown in Figure 5. The average effective number of Gaussian densities in the particle filter posterior density is 4, despite the particle filter using 2000 particles. Further analysis of the effective number of Gaussian densities revealed that the GPS bias estimates (10) augmented to the East, North, Heading state vector accounted for up to 5 additional mixture components needed to represent the posterior density. The posterior density is well represented by a relatively small term Gaussian mixture at each sample in the collected data. This compactness explains how the GSF with $N_g = 5$ terms in the posterior density after condensation is able to outperform the PF with 2000 particles.

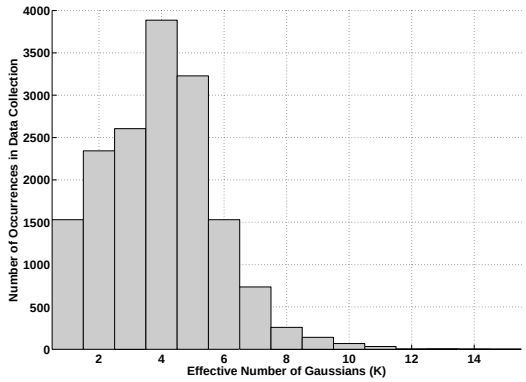

Fig. 5. The effective number of Gaussian densities in the particle filter data over the entire 37 minute data collection. The particle filter is using 2000 particles to represent a K-term Gaussian mixture with fewer than 6 components over much of the data collection.

VI. CONCLUSION

A Gaussian Sum filter was developed to perform localization using multi-modal vision measurements in the absence of GPS. The GSF uses a simple condensation technique to maintain the number of Gaussian components in the posterior density over time. The GSF remains converged with a precise global position estimate over an extended 32 minute GPS blackout. The GSF is successful in estimating the global position due to accurate representation of the multi-modal likelihood function that arises from the vision processing algorithms. The urban localization problem has posterior densities that are demonstrated to be well-represented with a small term Gaussian mixture, and adds to the representational appropriateness of the GSF. Data collected with the Cornell University autonomous vehicle shows the ability of the GSF to outperform the particle filter in terms of MMSE accuracy and representational appropriateness for the multi-modal localization problem.

VII. ACKNOWLEDGMENTS

The authors would like to thank Peter Moran and Aaron Nathan for assistance in collecting data and the members of the Cornell DARPA Urban Challenge Team for providing sensors, sensor interfaces, and other data collection support. This work is supported under the Northrop Grumman Electronic Systems Scholars Program.

REFERENCES

[1] E. Syed and M. Cannon, "Map-aided gps navigation," *GPS World*, vol. 16, no. 11, pp. 39–44, 2005.

[2] Y. Cui and S. S. Ge, "Autonomous vehicle positioning with gps in urban canyon environments," *Robotics and Automation, IEEE Transactions on*, vol. 19, no. 1, pp. 15–25, 2003, 1042-296X.

[3] C. Fouque and P. Bonnifait, "Tightly-coupled gis data in gnss fix computations with integrity testing," *International Journal of Intelligent Information and Database Systems*, vol. 2, no. 2, pp. 167–186, 2008.

[4] W. S. Wijesoma, K. W. Lee, and J. Ibanez-Guzman, "Motion constrained simultaneous localization and mapping in neighbourhood environments," in *Robotics and Automation, 2005. ICRA 2005. Proceedings of the 2005 IEEE International Conference on*, 2005, pp. 1085–1090.

[5] I. Miller and M. Campbell, "Particle filtering for map-aided localization in sparse gps environments," in *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on*, 2008, pp. 1834–1841.

[6] T. Bailey and H. Durrant-Whyte, "Simultaneous localization and mapping (slam): part ii," *Robotics and Automation Magazine, IEEE*, vol. 13, no. 3, pp. 108–117, 2006, 1070-9932.

[7] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking," *Signal Processing, IEEE Transactions on [see also Acoustics, Speech, and Signal Processing, IEEE Transactions on]*, vol. 50, no. 2, pp. 174–188, 2002.

[8] H. W. Sorenson and D. L. Alspach, "Recursive bayesian estimation using gaussian sums," *Automatica*, vol. 7, no. 4, pp. 465–479, 1971.

[9] D. Alspach and H. Sorenson, "Nonlinear bayesian estimation using gaussian sum approximations," *Automatic Control, IEEE Transactions on*, vol. 17, no. 4, pp. 439–448, 1972.

[10] D. J. Salmond, "Mixture reduction algorithms for uncertain tracking," Royal Aerospace Establishment, Tech. Rep. 88004, Jan. 1988.

[11] N. Peach, "Bearings-only tracking using a set of range-parameterised extended kalman filters," *Control Theory and Applications, IEE Proceedings* -, vol. 142, no. 1, pp. 73–80, 1995, 1350-2379.

[12] T. R. Kronhamn, "Bearings-only target motion analysis based on a multihypothesis kalman filter and adaptive ownship motion control," *Radar, Sonar and Navigation, IEE Proceedings* -, vol. 145, no. 4, pp. 247–252, 1998, 1350-2395.

[13] N. M. Kwok, G. Dissanayake, and Q. P. Ha, "Bearing-only slam using a sprt based gaussian sum filter," in *Robotics and Automation, 2005. ICRA 2005. Proceedings of the 2005 IEEE International Conference on*, 2005, pp. 1109–1114.

[14] I. Miller, M. Campbell, D. Huttenlocher, F.-R. Kline, A. Nathan, S. Lupashin, J. Catlin, B. Schimpf, P. Moran, N. Zych, E. Garcia, M. Kurdziel, and H. Fujishima, "Team cornell's skynet: Robust perception and planning in an urban environment," *Journal of Field Robotics*, vol. 25, no. 8, pp. 493–527, 2008.

[15] "Urban challenge rules," Defense Advanced Research Projects Agency (DARPA), Tech. Rep., 2007.

[16] B. D. O. Anderson and J. B. Moore, *Optimal filtering*, ser. Prentice-Hall information and system sciences series. Englewood Cliffs, N.J.: Prentice-Hall, 1979.

[17] Y. Bar-Shalom, X. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation*. New York: John Wiley and Sons, Inc, 2001.

[18] C. M. Bishop, *Pattern recognition and machine learning*. Springer-Verlag New York, Inc., 2006.

[19] J. H. Kotecha and P. M. Djuric, "Gaussian particle filtering," *Signal Processing, IEEE Transactions on [see also Acoustics, Speech, and Signal Processing, IEEE Transactions on]*, vol. 51, no. 10, pp. 2592–2601, 2003, 1053-587X.