Map-based Priors for Localization

Sang Min Oh, Sarah Tariq, Bruce N. Walker, and Frank Dellaert
College of Computing, Georgia Institute of Technology, Atlanta, GA

***Abstract*— Localization from sensor measurements is a fundamental task for navigation. Particle filters are among the most promising candidates to provide a robust and real-time solution to the localization problem. They instantiate the localization problem as a Bayesian filtering problem and approximate the posterior density over location by a weighted sample set. In this paper, we introduce map-based priors for localization, using the semantic information available in maps to bias the motion model toward areas of higher probability. We show that such priors, under a particular assumption , can easily be incorporated in the particle filter by means of a pseudo likelihood. The resulting filter is more reliable and more accurate. We show experimental results on a GPS-based outdoor people tracker that illustrate the approach and highlight its potential.**

I. Introduction

In this paper we introduce map-based priors for localization. Localization from sensor measurements is a fundamental task for navigation, but in many applications the available sensor readings are sometimes unreliable or not available altogether. As an example, at Georgia Tech we are designing a system to localize blind people in urban environments based on GPS, but it is well known that GPS is often unreliable in urban environments, e.g., due to satellite obstruction. A redeeming feature of urban environments, however, is that often high quality maps of the area are available, and hence the question arises whether we can use the information contained in these maps to aid in the tracking process. Our solution is to use a priori available maps to define a probability distribution over locations and use that to augment the motion model in a Bayesian filtering framework. The theory is valid for any Bayesian filtering framework, but the current paper is mostly concerned with its implementation within a particle filter, a sampling-based implementation of the Bayesian filter [1], [2].

A. Localizing Robots and People

There is an extensive literature on localizing both robots and people. A good overview of the robot localization and mapping literature can be found in [2]. Localization and tracking of people has been explored most thoroughly in the Augmented Reality (AR) community. While systems for indoors tracking have produced good results, outdoor tracking has proved to be a more challenging problem [3]. Most current applications do real-time outdoor tracking with a combination of GPS, dead reckoning and vision based techniques. GPS is not always reliable, since its accuracy depends crucially on the number of satellites visible. In urban areas the view of the satellites can be blocked by tall buildings or even foliage [4].

To deal with the problems of GPS or similar intermittently available sensors, many current systems fuse the information coming from multiple, complimentary sensors. Simple solutions use a cascade of sensors, with GPS usually being the primary one. [4] use GPS and compass information, [5] uses GPS and vision, [6] use GPS and dead reckoning. Others have proposed fusing multiple sensors using Extended Kalman filters [3], or fuse sensor information from vision and inertial trackers [7]. However, dead reckoning is prone to drift over time and can quickly accumulate considerable error. Vision-based sensors, on the other hand, typically need a large amount of data about the environment, and are not robust to lighting and abrupt motion changes.

B. Map-Based Priors for Localization

The main contribution of this paper is the introduction of maps as a prior in a Bayesian filtering paradigm. We use the semantic information available in maps to bias the motion model toward the areas of higher probability. The relative probabilities of different areas of the map reflect our beliefs about both the person and the environment. We may assign a higher probability to the edges of a sidewalk if we observe that blind people tend to walk more on the edges of a sidewalk than in its center. Similarly, blocked off roads would have higher probability for a pedestrian than regular roads, and paths that lead to centers of social interaction will be more probable than those leading to areas of low interest. Given no other information except such a probability map, a motion model, and the initial position and orientation of a person, we can estimate the short-term trajectory of this person by considering the most probable behavior. We show that such priors, under assumption of a particular mathematical form, can easily be incorporated in the particle filter by means of a pseudo likelihood.

C. Related Work

Maps have been used previously for robot localization. Laser range data [8], [9] or odometry [10] can be matched to an existing reference map to estimate location and orientation. For example, [1] have used maps of the environment which differentiate between obstacles and paths, but only for the process of culling useless particles. One observation is that sighted people generally walk more in the center of corridors. This can be realized by constraining the particles to move on a Voronoi diagram of the space [11].

D. Overview

The remainder of this paper is organized as follows: in the next section, Section II, we present the problem of localization as an instance of the Bayesian filtering problem and describe particle filters; Section III discusses map-based priors and a sampling technique with theoretical justification; Section IV contains experimental results illustrating the improved performance of our method over GPS alone.

II. Localization

A. The Bayes Filter

The localization problem can be expressed as a Bayesian filtering problem [1], [12]. Location is typically represented as a three-dimensional vector $\mathbf{x} = [x, y, \theta]^T$ encoding position and orientation, but can be more general, e.g., full 3D position and attitude with respect to a reference frame. In the Bayesian framework, we like to obtain the posterior density $P(x_t|Z_{1:t})$ over the current state x_t conditioned on all measurements $Z_{1:t} = \{z_i | i = 1..t\}$ up to time t. By Bayes rule this is expressed (up to a constant factor) as the product of a likelihood $P(z_t|x_t)$ and a prior $P(x_t|Z_{1:t-1})$ obtained recursively from measurements $Z_{1:t-1}$ up to time $t-1$:

$$P(x_t|Z_{1:t}) \propto P(z_t|x_t)P(x_t|Z_{1:t-1}) \tag{1}$$

We refer to $P(z_t|x_t)$ as the *measurement model* as it describes the probability of making observation z_t when the person is at location x_t. Thus, the measurement model is selected to capture the error characteristics of the sensor. The *predictive distribution* $P(x_t|Z_{1:t-1})$ denotes the probability of a person being in the location x_t at time t given the history of sensor measurements $Z_{1:t-1}$. We obtain the predictive distribution by integrating the *motion model* $P(x_t|x_{t-1})$ over the posterior $P(x_{t-1}|Z_{1:t-1})$:

$$P(x_t|Z_{1:t-1}) = \int_{x_{t-1}} P(x_t|x_{t-1})P(x_{t-1}|Z_{1:t-1}) \tag{2}$$

The motion model $P(x_t|x_{t-1})$ encodes the dynamics of the target as a conditional density of the current location x_t given x_{t-1}. Note that this motion model can be conditioned on additional information such as a control input u_t at time t, but for the sake of notational simplicity we leave this implicit below.

Combining (1) and (2) we obtain the *Bayes filter*,

$$P(x_t|Z_{1:t}) = k_t P(z_t|x_t) \int_{x_{t-1}} P(x_t|x_{t-1})P(x_{t-1}|Z_{1:t-1}) \tag{3}$$

where k_t is a normalizing factor. Thus, the posterior $P(x_t|Z_{1:t})$ over location is recursively obtained from the previous posterior $P(x_{t-1}|Z_{1:t-1})$, by integrating the predefined model of the target dynamics and the sensor measurements.

B. Particle Filters

Particle filters [1], [13], [14] take an importance sampling approach to implement the Bayes filter (3). They approximate the posterior $P(x_{t-1}|Z_{1:t-1})$ as a weighted sample set,

$$P(x_{t-1}|Z_{1:t-1}) \approx \sum_{i=1}^{N} w_{t-1}^{(i)} \delta(x_{t-1}^{(i)}, x_{t-1}) \tag{4}$$

where location $x_{t-1}^{(i)}$ and weight $w_{t-1}^{(i)}$ are the information stored in ith particle, and N is the number of samples. The importance sampling approach is applied to sample efficiently from the posterior density (3) [14]. In the importance sampling approach, we sample from the *proposal distribution* and we weight each sample with the relevant weight. In the case of the localization, the proposal distribution is the predictive density (2) and the individual weight is obtained from the measurement model $P(z_t|x_t)$. Thus, the first step to estimate the current posterior density $P(x_t|Z_{1:t})$ recursively from the previous posterior $P(x_{t-1}|Z_{1:t-1})$ is to compute the predictive distribution from which we can efficiently sample. With the given representation, we can approximate the empirical predictive distribution (2), i.e., the mixture model :

$$P(x_t|Z_{1:t-1}) \approx \sum_{i=1}^{N} w_{t-1}^{(i)} P(x_t|x_{t-1}^{(i)}) \tag{5}$$

Above, the mixture coefficients are the sample weights $w_{t-1}^{(i)}$, and the mixture component $P(x_t|x_{t-1}^{(i)})$ is the motion model for each sample $x_{t-1}^{(i)}$. This empirical predictive distribution is then used as the proposal distribution $Q(x_t)$ for the importance sampling, from which we obtain unweighted samples $\hat{x}_t^{(j)}$:

$$\hat{x}_t^{(j)} \sim Q(x_t) = \sum_{i=1}^{N} w_{t-1}^{(i)} p(x_t|x_{t-1}^{(i)}) \tag{6}$$

To sample from (6), one first chooses a component i at random, according to the weights $w_{t-1}^{(i)}$, and then sample from the corresponding component $P(x_t|x_{t-1}^{(i)})$. This is done N' times, where N' can equal to N, or can be adapted to the complexity of the hypothesis or available processing power [15]. The unweighted samples $\{\hat{x}_t^{(j)}\}_{j=1}^{N'}$ are then upgraded to the current posterior $P(x_t|Z_{1:t})$, yielding the importance weights $w_t^{(j)}$:

$$w_t^{(j)} = \frac{k_t P(z_t|\hat{x}_t^{(j)}) P(\hat{x}_t^{(j)}|Z_{1:t-1})}{Q(\hat{x}_t^{(j)})} \approx k_t P(z_t|\hat{x}_t^{(j)}) \tag{7}$$

Thus, the current posterior $P(x_t|Z_{1:t})$ is approximated by the following newly obtained weighted sample set :

$$P(x_t|Z_{1:t}) \approx \sum_{j=1}^{N'} w_t^{(j)} \delta(x_t^{(j)}, x_t) \tag{8}$$

The key advantage of particle filters is that they can represent arbitrary posterior probability distributions, and can deal with arbitrarily complex measurement models.

III. MAP-BASED PRIORS

A. Localization with Map-based Priors

When localizing either a robot or a person in a known environment, it would be beneficial to be able to use the available semantic information in maps to bias the motion model $P(x_t|x_{t-1})$ toward the areas of higher probability. The map in figure 1 shows the probability of being in a region with colors. The zone with the brighter color is the area with the high probability while the zone with the darker color is the area with the low probability. The black zones denote the zero probability areas which may include the buildings and shrubs if we consider only outdoor localization. Note that the map can either contain information about the outdoor or indoor environments or both. Denoting the map by M, the Bayes filter (3) in this case now depends on M :

$$P(x_t|Z_{1:t},M) = k_t P(z_t|x_t) \int_{x_{t-1}} P(x_t|x_{t-1},M) P(x_{t-1}|Z_{1:t-1},M) \tag{9}$$

where now the *augmented motion model* $P(x_t|x_{t-1},M)$ is conditioned on the information in the map M.

The use of pre-existing maps generates a more informed posterior $P(x_t|Z_{1:t},M)$ by exploiting knowledge about the environment. For example, people tend to walk on the sidewalks rather than on the road or on the grass, robots tend to stay away from objects due to on-line obstacle avoidance methods, etc.

However, there are two potential difficulties :

1) While it is feasible to either hand-build or learn a map-based prior $P(x_t|M)$ over locations x_t, it is not immediately obvious how to combine this information with the unconditional motion model $P(x_t|x_{t-1})$ to obtain the augmented motion model $P(x_t|x_{t-1},M)$.
2) Depending on the nature of the augmented motion model $P(x_t|x_{t-1},M)$, evaluating the integral in (9) is potentially much more difficult.

Below we show that both difficulties are overcome using a particular form for the augmented motion model.

B. Augmented Motion Model

A map-based prior $P(x_t|M)$ and the unconditional motion model $P(x_t|x_{t-1})$ can be combined in an augmented motion model $P(x_t|x_{t-1},M)$ if we require the augmented motion model to be applied to *local transitions*. The local transition assumption over the augmented motion model states that the transition from the previous location x_{t-1} to the current location x_t occurs only locally, i.e. the model considers only the short transitions between two close locations and the transition is not influenced by the global map structures. Under this assumption, all possible transitions can be classified into two cases :

1) The current location x_t and the previous location x_{t-1} are in the same probability zone. Thus, the map-based priors for the two locations are equal, $P(x_t|M) = P(x_{t-1}|M)$.
2) The location x_t and x_{t-1} belong to the distinct probability zones. In this case, the map-based priors for the locations differ, $P(x_t|M) \neq P(x_{t-1}|M)$.

In the first case the augmented motion model $P(x_t|x_{t-1},M)$ is the same as the unconditional motion model $P(x_t|x_{t-1})$. This is because the local transition assumption will allow the transition to be unaware of the global map M.

However, in the second case of the inter-zone transition, the augmented motion model $P(x_t|x_{t-1},M)$ should be adjusted by the relative ratio of the map-based priors between the two locations $\frac{P(x_t|M)}{P(x_{t-1}|M)}$. This results in an augmented motion model which is biased toward the the areas of higher probability. In general, the inside-zone transition case is a special case of the inter-zone transition case with the map-based prior ratio $\frac{P(x_t|M)}{P(x_{t-1}|M)}$ equal to one. Thus, under the particular assumption, we obtain the augmented motion model :

$$P(x_t|x_{t-1},M) = \alpha_t P(x_t|x_{t-1}) (\frac{P(x_t|M)}{P(x_{t-1}|M)})^{\beta} \tag{10}$$

$$= \alpha'_t P(x_t|x_{t-1}) P(x_t|M)^{\beta} \tag{11}$$

where the α'_t is a normalizing constant and β denotes the relative importance of the map-based prior $P(x_t|M)$ with respect to the unconditional motion model $P(x_t|x_{t-1})$.

There is another justification for the resulting augmented motion model (11), based on the Gibbs distribution $E(x)$:

$$P(x) \triangleq e^{-E(x)} \tag{12}$$

The Gibbs distribution $E(x)$ is interpreted as an energy function of $P(x)$. The x that maximizes $P(x)$ equals the x that minimizes the energy function $E(x)$. Thus, the augmented motion model can be obtained by defining its energy function $E(x)$ which can also be interpreted as the penalty function. Natural phenomenon prefers lower energy states and accordingly there is less penalty in such states. Thus, the energy function of the augmented motion model $P(x_t|x_{t-1},M)$ represents the penalty at the point x_t around the specific location x_{t-1} given the map M. Again, if we focus on the local transitions only, we can claim that the penalty at a location x_t is the linear sum of the penalty which is unconditioned on map M and the penalty that results from the inter-zone transition.

The assumption of *additive penalty* results in the following augmented motion model :

$$P(x_t|x_{t-1},M) = \alpha'_t \exp\left\{\log P(x_t|x_{t-1}) + \beta' \log P(x_t|M)\right\} \tag{13}$$

Above, α' is a normalizing constant and β' is a parameter that balances the map-based prior $P(x_t|M)$ the unconditional motion model $P(x_t|x_{t-1})$. The assumption of the additive penalty model is intuitively appealing, and the resulting model in (13) equals the model (11) exactly.

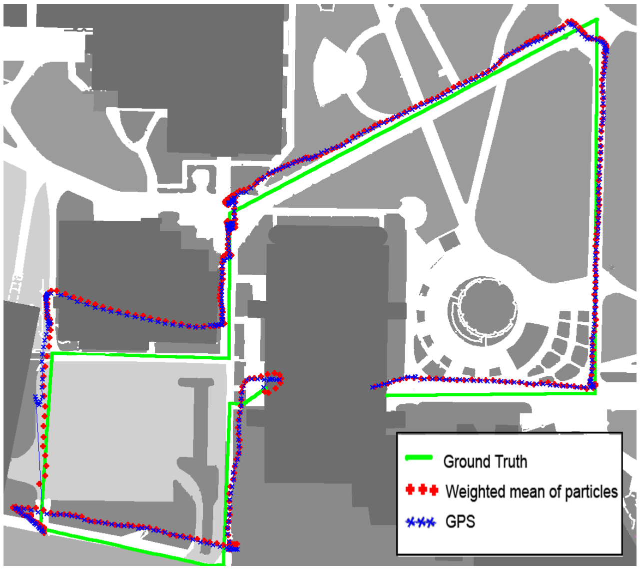

Fig. 1. Tracking without using a map prior on the first dataset. Note that the particle trace is almost coincident with the GPS trace.

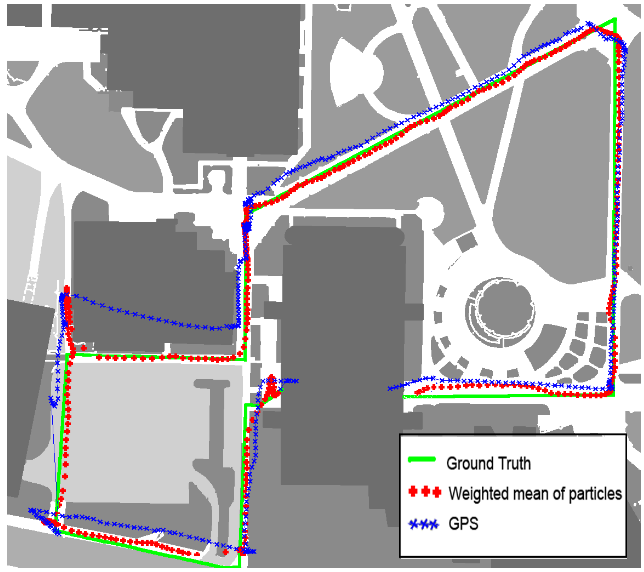

Fig. 2. Tracking using a map-based prior on the first dataset. Compared with Figure 1, the tracker now prefers to stay on the heavily traveled areas of the map such as sidewalks (white in the map). The resulting path is much closer to the ground-truth path, as a result of mediating the noisy GPS data with map-based prior.

C. Incorporating Map Priors in the Bayes Filter

Plugging the new motion model (11) into the augmented Bayes filter (9), the map prior factor $P(x_t|M)^\beta$ can be moved out of the integral, as it does not depend on x_{t-1}:

$$P(x_t|Z_{1:t},M) = k''_t P(z_t|x_t)P(x_t|M)^\beta P'(x_t|Z_{1:t-1},M) \quad (14)$$

where

$$P'(x_t|Z_{1:t-1},M) = \int_{x_{t-1}} P(x_t|x_{t-1})P(x_{t-1}|Z_{1:t-1},M)$$

is the predictive density without taking the map prior into account . This has the same form as Equation 3 on page 2, but now with an additional factor $P(x_t|M)^\beta$ that mediates the information given by the map prior $P(x_t|M)$.

The implication of the particularly simple form of the augmented filter (14) is that incorporating the map prior in the particle filter is straightforward. We can propose samples from the same proposal density $Q(x_t)$ (6), and simply augment the likelihood weight with a factor $P(x_t|M)^\beta$ derived from the map prior:

$$w_t^{(j)} = \frac{k''_t P(z_t|x_t^{(j)})P(x_t^{(j)}|M)^\beta P'(x_t^{(j)}|Z_{1:t-1},M)}{Q(\hat{x}_t^{(j)})} \quad (15)$$

$$= k''_t P(z_t|\hat{x}_t^{(j)})P(\hat{x}_t^{(j)}|M)^\beta \quad (16)$$

In other words, we have exactly the same particle filter as before, but now we multiply each importance weight by an additional map derived factor $P(x_t|M)^\beta$.

IV. Results

A. Experimental Setup

We show results on two different datasets collected with a GPS receiver in outdoor environments. The Global Positioning System (GPS) is widely known as a solution for outdoors localization. The GPS receiver used to gather the data is a handheld Garmin GPS 72s receiver that has an error of less than 15 meters 95% of the time. Data from GPS is converted from Latitude/Longitude coordinates to map coordinates using the assumption that the curvature of the Earth is negligible over the small region of interest. Two GPS datasets were gathered on the Georgia Tech campus, and used as the test data for the localization with/without map-based priors. In addition, we assume a Gaussian measurement model for the GPS.

B. Building the Map Prior

The areas in the maps fall into one of the six possible zones and are shown in different colors as in Figure 1. The areas, corresponding colors and the assigned probabilities are as follows in order of decreasing probability : sidewalks (white, 60%), parking lots and blocked off roads (light gray, 30%), lawns (gray, 5%), roads (dark gray, 5%), shrubs (very dark gray, 0%) and buildings (black, 0%). The ratio of two probabilities determines how easily a particle can cross the boundary of two areas. The relative importance ratio β in Equation 11 is set to be 1.0 in all the experiments.

C. Data set 1

Figures 1 and 2 show that results obtained using the map-based prior approach are more accurate, stable and keep closer to the actual path most of the time. The blue crosses are the trace of the actual GPS readings, the green line is the ground truth, and the red crosses are the weighted mean of the particle cloud at each time step. In figure 1, we show the behavior of the particle filter *without* using a map prior. Due to erroneous GPS readings the particles pass through the shrubs at the bottom left and pass through the building at the left center. In addition, the particles pass through the lawns during most of the path in the upper right. By using the map-based priors we can see in 2 that the particles avoid passing through the shrubs and buildings, and stay on the center of the sidewalks most of the time. The advantage of the map priors is most clearly evident when the GPS readings are noisy and unstable. This tends to happen around tall buildings, where the satellite reception gets bad. On the left side of figure 2, we can see that the particles correctly follow the sidewalk where the GPS readings erroneously indicate that the person is going through a building or shrubs. In this case the particles tend to follow the sidewalk, since it is the highest probability area within the GPS error radius.

# \ var	5 m	10 m	15 m
500	3.01 / 2.43 / 20	2.96 / 2.37 / 20	2.96 / 2.48 / 16
1000	3.03 / 2.42 / 20	2.98 / 2.41 / 19	2.98 / 2.51 / 16
1500	3.03 / 2.43 / 20	2.98 / 2.44 / 18	2.98 / 2.46 / 17

TABLE I

RMSE RESULTS FOR THE FIRST DATASET. THE ROWS ARE THE NUMBER OF PARTICLES USED AND THE COLUMNS ARE THE ASSIGNED GPS VARIANCES. EACH CELL SHOWS THREE ESTIMATES : THE RMSE OF THE NO-MAP VERSION LOCALIZATION IN METER, MAP-PRIOR BASED LOCALIZATION IN METER AND THE PERCENTAGE OF IMPROVEMENT BY USING MAP-PRIORS.

To obtain a more quantitative evaluation of our method, we systematically varied several key parameters and tabulated the results. Table I shows the root mean squared errors (RMSE) of each experiment along with different values for the number of particles and GPS variance. The RMSE gives a measure of how much the estimated path differs from the ground truth :

$$RMSE = \sqrt{\frac{\sum_{i=1}^{n}(l_i - t_i)^2}{n}} \quad (17)$$

RMSE in equation 17 is calculated by summing the squared distances of each estimated location l_i from the ground truth t_i, which is estimated by interpolating between known waypoint locations, and dividing by the total number of data points n.

It is evident that the map-based priors perform better when the error of measurement model is set to be 10 meters or less. This is expected since the Garmin GPS receiver is known to have an error <15 meters 95% of the time, which under our Gaussian measurement model assumption translates to a variance of approximately 7.6 meters. Moreover, a larger number of particles does not seem to bring additional positive effect; the RMSE of map-prior based localization is relatively static for a constant GPS variance. This is because 500 particles can efficiently approximate the posterior in Equation 8 without need for additional number of particles which leads to more processing requirements. This is a particularly encouraging result for a mobile application where processing resources may be limited. For the results shown in figures 1 and 2 the number of particles are set to be 500 and the GPS variance is set to be 5 meters.

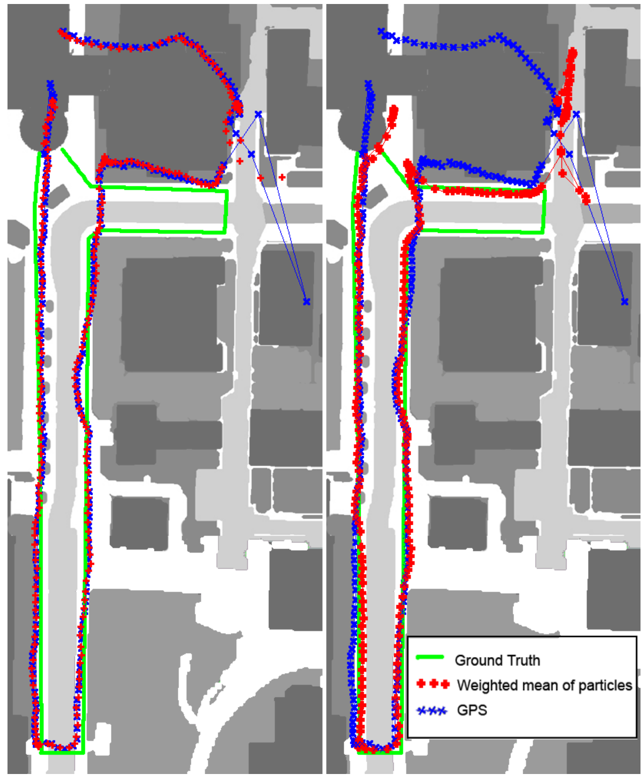

Fig. 3. (a) Left. Tracking without map priors.
(b) Right. Tracking results with map priors.

D. Data set 2

The second dataset shows a similar improvement in tracking due to the use of map-based priors. Figure 3(a) shows the tracking result without map-based priors and the result in figure 3(b) is obtained by applying the map-based priors. On the top right of 3(a) we note that, similar to the first dataset, the particles follow the GPS through the buildings, while the particles in 3(b) stays in the higher probability area nearby. Similarly, the particles in 3(a) go into shrubs and the road at the bottom of the map due to the erroneous GPS readings, but the particles in 3(b) constantly stay near the center of the sidewalks. Both trajectories deviate from the true green path at the top right of the map, due to the extremely bad GPS readings. The number of particles are set to be 500 and the GPS variance is set to be 5 meters for the results in figures 3(a) and 3(b).

The RMSE results with different parameter settings are summarized in table II. As in table I, the results show the overall advantage of our map-based priors approach.

Even though the advantage of using the map-based priors is evident in the figures, the RMSE does not always reflect

# \ var	5 m	10 m	15 m
500	4.21 / 3.53 / 16	4.02 / 3.43 / 15	3.85 / 3.47 / 10
1000	4.06 / 3.56 / 12	4.01 / 3.42 / 15	3.84 / 3.59 / 7
1500	4.09 / 3.56 / 13	4.03 / 3.39 / 16	3.88 / 3.69 / 5

TABLE II

RMSE RESULTS FOR THE SECOND DATASET.

this. For example, one of the results (using 1500 particles and 15 meters GPS variance) in table II shows only a 5 percent performance improvement by using map-priors. There are two reasons for this. Firstly, the errors between the estimated and the true path are small for most of the path in both datasets, since the GPS tends to follow the ground truth closely in open areas. This outweighs the larger errors in some sections of the path, and brings the overall RMSE for the two trajectories, which are generated with/without map-based priors, closer. Secondly, the RMSE calculates error based solely on the distance between the estimated and true locations and does not take into account the context of the estimated locations. For example, at the top right part of figure 3(a) the estimates lie inside a building, and RMSE fails to assign higher error for those points than to estimates that lie the same distance away from the ground truth but on the sidewalk, as in the top right of figure 3(b). Hence, even though the results in figure 3(b) are visually better, there is only a slight decrease in RMSE.

In summary, the map-based priors provide an effective and efficient means of improving the accuracy of localization schemes. We showed the improved tracking results qualitatively in figures 2 and 3(b), and quantitatively in tables I and II. Specifically, map-prior based localization produces more logical results in the presence of large sensor noise and brings the estimated path closer to the ground truth.

V. Conclusions and Future Work

We introduced map-based priors, with which we can efficiently perform localization using particle filters even when the main sensor is inaccurate or unreliable at times. As the experimental results show, even the localization using noisy sensors results in far more stable local tracking, representing the ground truth route more correctly. This technique is applicable to variety of applications, including tracking robots and humans.

Even though the initial results are promising, there is plenty to do in terms of future work. In fact, the system we have implemented can be considered a bare-bones version, so improvements in most areas should lead to even better results. For example, the accuracy of our method relies on the quality of the probability map. A more realistic and useful map could be obtained by taking observations of the area in question over a period of time. Furthermore, a time-dependent map could be built in this way. Such maps should lead to improvements in the accuracy of the method. Also, in our particular application, the results we show are not as good as they could be if we had used a better error model for GPS. In particular, the Gaussian model we used is not very appropriate to model the systematic errors associated with GPS.

As a future application of the proposed localization technique, we plan on integrating the algorithm into the Georgia Tech System for Wearable Audio Navigation (SWAN), a mobility tool for the visually impaired. The SWAN system will provide guidance both outdoors and indoors, with GPS serving as the primary sensor outdoors and computer vision as the primary sensor indoors. In conjunction with inertial sensors and the map-based priors discussed in this paper, we expect an effective and successful practical application, of benefit to a great number of people.

Acknowledgments

The authors would like to acknowledge graduate students Lars Fiedler and Jeffrey Lindsay for their assistance in gathering real-time GPS data with the SWAN system.

References

[1] F. Dellaert, D. Fox, W. Burgard, and S. Thrun, "Monte Carlo Localization for mobile robots," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 1999.

[2] S. Thrun, D. Fox, F. Dellaert, and W. Burgard, "Particle filters for mobile robot localization," in *Sequential Monte Carlo Methods in Practice* (A. Doucet, N. de Freitas, and N. Gordon, eds.), New York: Springer-Verlag, January 2001.

[3] R. Azuma, B. Yohan, R. Behringer, S. Feiner, and S. J. B.MacIntyre, "Recent advances in augmented reality," *Computers and Graphics*, November 2001.

[4] S. Feiner, B. MacIntyre, T. Hollerer, and A. Webster, "A touring machine: Prototyping 3D mobile augmented reality systems for exploring the urban enviroment," in *Proc of ISWC*, pp. 74–81, 1997.

[5] R. Behringer, "Registration for outdoor augmented reality applications using computer vision techniques and hybrid sensors," in *Proceedings of IEEE Virtual Reality*, pp. 244–251, March 1999.

[6] A. Helal, S. Moore, and B. Ramachandran, "Drishti: An Intergrated Navigation System for Visually Impaired and Disabled," in *Proceedings of the 5th International Symposium on Wearable Computer*, (Zurich,Switzerland), October 2000.

[7] Y. Yokakohiji, "Accurate image overlay on video see-through HMDs using vision and accelerometers," in *IEEE VR*, pp. 247–254, 2000.

[8] J. Borenstein, B. Everett, and L. Feng, *Navigating Mobile Robots: Systems and Techniques*. Wellesley, MA: A. K. Peters, Ltd., 1996.

[9] G. Borges and M.J.Aldon, "An optimal pose estimator for map-based mobile robot dynamixc localization: Experimental comparison with the EKF," in *IEEE International Conference on Robotics and Automation (ICRA)*, May 2001.

[10] N. Ayache and O. Faugeras, "Maintaining representations of the environment of a mobile robot," *IEEE Trans. on Robotics and Automation*, vol. 5, no. 6, pp. 804–819, 1989.

[11] L.Liao, D.Fox, J.Hightower, H.Kautz, and D.Schulz, "Voronoi tracking: Location estimation using sparse and noisy sensor data," in *IROS*, 2003.

[12] D. Fox, J. Hightower, L. Liao, D. Schulz, and G. Borriello, "Bayesian filtering for location estimation," *IEEE Pervasive Computing*, pp. 10–19, July 2003.

[13] M. Isard and A. Blake, "Contour tracking by stochastic propagation of conditional density," in *Eur. Conf. on Computer Vision (ECCV)*, pp. 343–356, 1996.

[14] S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for on-line non-linear/non-Gaussian Bayesian tracking," *IEEE Transactions on Signal Processing*, vol. 50, pp. 174–188, Feb. 2002.

[15] C. Kowk, D. Fox, and M. Meilă, "Adaptive real-time particle filters for robot localization," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2003.