Detection/Tracking of moving targets with synthetic aperture radars

Gregory E. Newstadt[a], Edmund Zelnio[b], LeRoy Gorham[b], and Alfred O. Hero III[a]

[a]Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109, USA;

[b]Air Force Research Laboratory, Wright Patterson Air Force Base, OH, 45433, USA

ABSTRACT

In this work, the problem of detecting and tracking targets with synthetic aperture radars is considered. A novel approach in which prior knowledge on target motion is assumed to be known for small patches within the field of view. Probability densities are derived as priors on the moving target signature within backprojected SAR images, based on the work of Jao.[1] Furthermore, detection and tracking algorithms are presented to take advantage of the derived prior densities. It was found that pure detection suffered from a high false alarm rate as the number of targets in the scene increased. Thus, tracking algorithms were implemented through a particle filter based on the Joint Multi-Target Probability Density (JMPD) particle filter[2] and the unscented Kalman filter (UKF)[3] that could be used in a track-before-detect scenario. It was found that the PF was superior than the UKF, and was able to track 5 targets at 0.1 second intervals with a tracking error of $0.20 \pm 1.61m$ (95% confidence interval).

Keywords: Synthetic Aperture Radar, target tracking, target detection, JMPD, particle filter

1. INTRODUCTION

The ability to track moving targets with airborne radar is a problem that has drawn considerable interest from both the academic and government communities. Moreover, in many cases the opportunity now exists for continuous observation of regions of interest with airborne radars. Thus, there is great incentive to be able to efficiently detect and track individual vehicles over large areas.

Currently, radar systems operate in several modes depending on the application. In Moving Target Indication (MTI) mode, the radar focuses a narrow beam over small regions in the field of view for small integration times on the order of milliseconds. However MTI-mode radars suffer from tradeoffs between the small integration time (that leads to low SNR values and coarse resolutions) as compared to the number of regions that can be observed. Furthermore, MTI systems can only detect velocities in a single direction, which allow for evasive maneuvers to avoid detection.

In synthetic aperture radar (SAR) mode, an image is constructed by integrating radar pulses from spatially diverse points in the platform's trajectory. This 'synthetic aperture' leads to 2-dimensional imaging as well as much finer resolutions than in MTI-mode, due to the ability to use longer integration times. However, all of these benefits do not come for free. SAR was designed to image stationary scenes. Consequently, moving targets cause phase errors in the reconstruction of a SAR image that lead to smearing and displacement of a target's energy[1,4] as shown in Figure 1 using the toolbox provided by Gorham and Moore.[5]

Nevertheless, there have been attempts to use SAR for moving target detection. Werness[6] and Kirscht[7] propose methods based on estimating these phase errors, and then refocusing the images. Fienup[4] showed a sharpness ratio could be used to reduce the number of false alarms. However, in all of these cases, a large computational burden was involved in the processing of an image, limiting the field of view (FOV) that could be tracked.

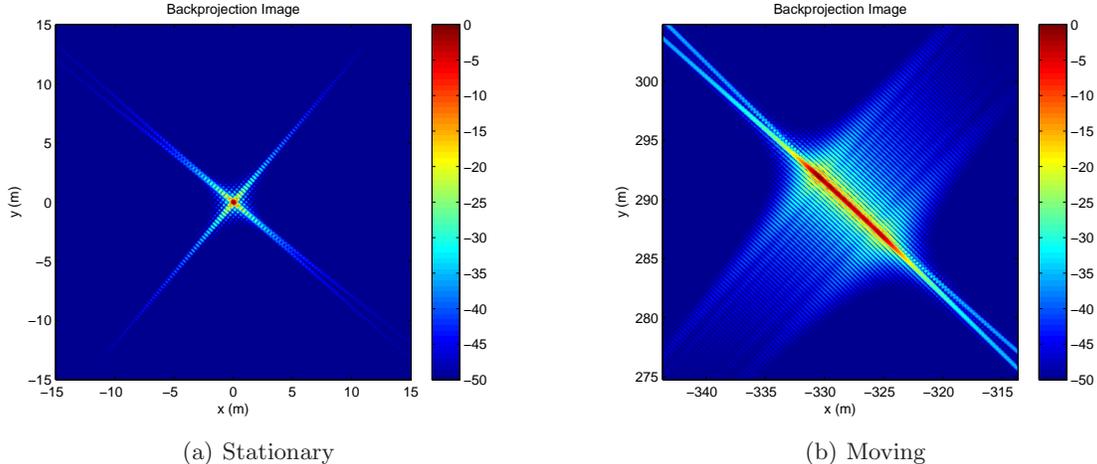(a) Stationary   (b) Moving

Figure 1. Here SAR images constructed through the backprojection method provided by Gorham and Moore[5] are shown for point targets. In (a) the point target is stationary at $(0,0)$ and the majority of the energy is focused at that point. In (b) the point target has velocity $(v_x, v_y) = (30,5)$ m/s and acceleration $(a_x, a_y) = (3,1)$ m/s^2. The target is both displaced in the image (by more than 300 meters) and smeared (with smear length of about 10 meters).

To combat the computational problem, a new approach is considered where road structures are used as prior knowledge on the trajectories of targets within the field of view. Jao[1] showed that the target's signature within the reconstructed image can be predicted geometrically given both the platform and the target's trajectories. In practice, the platform trajectory is supplied by onboard GPS and IMU data, and it can be assumed that it is known deterministically. In this work, it is assumed that the target's trajectory can be modeled based on known road structures, time of day, and traffic patterns. For the purposes of this paper, this model is assumed to be known a priori, though future work will tackle the problem of learning and/or refining the model based on received data.

The first contribution of this work is the derivation of prior densities for the location of a moving target's signature within a reconstructed SAR image. The second contribution of this paper is the development of a detection algorithm that takes advantage of these priors. For a single target, the algorithm was found to have reasonably good performance (75% probability of detection at 20% false alarm rate) over a simulated data set. However, when multiple targets were present within neighboring cells, the detection performance suffered (see Figure 5 for details). The third contribution of this paper helps combat this issue by providing two tracking algorithm based on (1) the joint multitarget probability density particle filter[2] (PF) and (2) the unscented Kalman filter[3] (UKF). It was found that PF was superior, with the ability to track 5 targets at 0.1 second intervals with a tracking error of 0.20 ± 1.61m over a field of view (FOV) of size 300x450 m^2 (as compared to an average tracking error of 5.67 ± 12.88m for the UKF over the same FOV).

The rest of the paper is organized as follows. In section 2 we formally introduce the problem statement. In section 3, the theoretical results are presented. Section 4 provides performance analysis based on a real dataset for both the detection and tracking algorithms. Finally, we conclude and point to future work in section 5.

2. PROBLEM STATEMENT

This section provides the basics of the SAR image formation process, using the same notation as Jao and the target trajectory model that will be used in the subsequent derivations.

For a 95% confidence interval

2.1 SAR image formation

For simplicity, it is assumed that all targets are point scatterers in a monostatic radar, with continuous back-projection. From Jao,[1] the signal sample g can be written as

$$g(t, \tau, \mathbf{r}(\tau)) = \frac{b(\tau, \mathbf{r}(\tau))}{R^2(\tau, \mathbf{r}(\tau))} \exp\left[j2\pi\nu_0 \left(t - \frac{2}{c}R(\tau, \mathbf{r}(\tau))\right)\right] p_r\left(t - \frac{2}{c}R(\tau, \mathbf{r}(\tau))\right) \tag{1}$$

where $\mathbf{r}(\tau) = (r_x, r_y, r_z)$ is the true position of a point scatterer, t is the fast-time, corresponding to the slant range $ct/2$, $\tau \in [T_0, T_1]$ is the slow-time, $R(t, \tau) = \|\mathbf{r}(\tau) - \mathbf{q}(\tau)\|_2$ is the range from platform to the point scatterer, c is the speed of light, $p_r(\cdot)$ is the output of the range compression filter, ν_0 is the radar center frequency, and $b(\tau, \mathbf{r}(\tau))$ is an amplitude variable that lumps together all effects other than spherical propagation such as backscatter, antenna gain, etc.. Furthermore, let $\mathbf{p} = (p_x, p_y, p_z)$ be a point in our reconstructed image (usually p_z is equal to zero or a function of p_x and p_y if a DEM is given) and $\mathbf{q}(\tau) = (q_x, q_y, q_z)$ be the true position of the platform. Then the image response at $\mathbf{p}$, $f(\mathbf{p})$, is obtained by integrating g over $\tau \in [T_0, T_1]$ and compensating for the time delay $(2/c)R(\tau, \mathbf{p})$. Normalizing to have a maximum unity amplitude, f can be written as

$$f(\mathbf{p}) = C \int_{T_0}^{T_1} \frac{b(\tau, \mathbf{r}(\tau))}{R^2(\tau, \mathbf{r}(\tau))} \exp\left[j2k_0 \left(\|\mathbf{p} - \mathbf{q}(\tau)\|_2 - \|\mathbf{r}(\tau) - \mathbf{q}(\tau)\|_2\right)\right]$$

$$\times p_r\left(\frac{2}{c}\left(\|\mathbf{p} - \mathbf{q}(\tau)\|_2 - \|\mathbf{r}(\tau) - \mathbf{q}(\tau)\|_2\right)\right) d\tau \tag{2}$$

When the point scatterer is stationary (i.e., $\mathbf{r}(\tau) = \mathbf{r}_0$ for all τ), then the maximum of equation (2) occurs at $\mathbf{p} = \mathbf{r}_0$. However, for general $\mathbf{r}(\tau)$, Jao shows that the local maximum of equation (2) occurs at $\mathbf{p} = \mathbf{p}^*$ when the following two conditions are met:

1. The Doppler shifts are equivalent, leading to

$$\left[\frac{d}{d\tau}\left(\|\mathbf{p} - \mathbf{q}(\tau)\|_2 - \|\mathbf{r}(\tau) - \mathbf{q}(\tau)\|_2\right)\right]_{\mathbf{p}=\mathbf{p}^*} = 0 \tag{3}$$

   This can be equivalently stated as

$$-\frac{\dot{\mathbf{q}}(\tau) \cdot [\mathbf{p}^* - \mathbf{q}(\tau)]}{\|\mathbf{p}^* - \mathbf{q}(\tau)\|_2} = \frac{[\dot{\mathbf{r}}(\tau) - \dot{\mathbf{q}}(\tau)] \cdot [\mathbf{r}(\tau) - \mathbf{q}(\tau)]}{\|\mathbf{r}(\tau) - \mathbf{q}(\tau)\|_2} \tag{4}$$

2. The ranges are equivalent, leading to

$$\|\mathbf{p}^* - \mathbf{q}(\tau)\|_2 = \|\mathbf{r}(\tau) - \mathbf{q}(\tau)\|_2 \tag{5}$$

When $\mathbf{r}(\tau)$, $\dot{\mathbf{r}}(\tau)$, $\mathbf{q}(\tau)$, and $\dot{\mathbf{q}}(\tau)$ are known, and $p_z = 0$, these equations can be solved explicitly for (p_x, p_y). To conserve space, the exact expressions are provided in a technical report to be released a later time. When the range compression filter is ideal, (i.e., p_r is just a delta function), and ignoring the amplitude variable, the image model is

$$f_{ideal}(\mathbf{p}) = \int_{T_0}^{T_1} \delta(\mathbf{p} - \mathbf{p}^*(\tau))d\tau \tag{6}$$

where $\mathbf{p}^*(\tau)$ satisfies the above constraints at each impulse $\tau \in [T_0, T_1]$. In practice, the image is constructed through a discrete set of impulses in a set $\tau \in T$ (with endpoints $T_0, T_1 \in T$) and the image will be corrupted by noise. Thus, a corrected model for the image is

$$f(\mathbf{p}) = I_{targ} \sum_{\tau \in T} \delta[\mathbf{p} - \mathbf{p}^*(\tau)] + \nu(\mathbf{p}) \tag{7}$$

where I_{targ} is an indicator variable for the existence of a target within the image, and $\nu(\mathbf{p})$ is i.i.d. noise. In this paper, it is assumed that $\nu(\mathbf{p}) \sim \mathcal{N}(0, \sigma_v^2)$ is zero-mean and Gaussian.
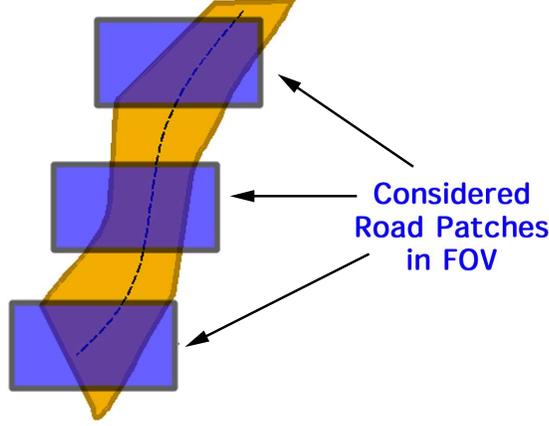
Figure 2. The strategy for detecting targets is illustrated here. Patches within the field of view are scanned along known road structures. Each road patch is associated with one or more target motion models. For example, near an intersection one may consider separate motion models for both stopped vehicles versus accelerating vehicles.

2.2 Target trajectory model

In practice only the platform trajectory $(\mathbf{q}(\tau), \dot{\mathbf{q}}(\tau))$ is known (through GPS and IMU data onboard the aircraft). However, it is assumed in this paper that we can model the target's motion within small regions of the FOV by taking advantage of road structures, traffic models, etc.. This leads to a target trajectory model:

$$\mathbf{r}(\tau) \quad = \quad \mathbf{r}_0 + \mathbf{v}_r \tau + \mathbf{a}_r \tau^2 / 2 \tag{8}$$

where $\mathbf{r}(\tau) = [x(\tau), \ y(\tau)]^T$, $\mathbf{r}_0 = [x_0, \ y_0]^T$, $\mathbf{v}_r = [v_x, \ v_y]^T$, and $\mathbf{a}_r = [a_x, \ a_y]^T$ are all mutually independent. Moreover, these quantities are modeled stochastically as

$$\mathbf{r}_0 \quad \sim \quad Uniform([x_{min}, \ x_{max}] \times [y_{min}, \ y_{max}]) \tag{9}$$

$$\mathbf{v}_r \quad \sim \quad \mathcal{N} \left(\mu_{\mathbf{v}}, \begin{bmatrix} \sigma_{vx}^2 & 0 \\ 0 & \sigma_{vy}^2 \end{bmatrix} \right) \tag{10}$$

$$\mathbf{a}_r \quad \sim \quad \mathcal{N} \left(\mu_{\mathbf{a}}, \begin{bmatrix} \sigma_{ax}^2 & 0 \\ 0 & \sigma_{ay}^2 \end{bmatrix} \right) \tag{11}$$

This model implicitly assumes a unique approach for moving target detection as compared to the existing methods; rather than searching for targets anywhere in the FOV, this strategy employs a 'zone defense' by searching for targets over small patches in the FOV. It is inherently assumed that the above model is accurate for small integration times (on the order of 1 second). Figure 2 illustrates this strategy for three consecutive patches along a single road. Note that it is possible for a single road patch to contain more than one considered target trajectory model. This may be necessary in situations where the target may exhibit several types of behavior (for example, acceleration/deceleration at an intersection).

It is worthwhile to note that for a detection strategy, one must be able to account for uncertainty in all of the states of the target (i.e., $(\mathbf{r}, \mathbf{v}, \mathbf{a})$). However, for the tracking algorithms, some of these variables are known from previous states, and updated through a state-transition model. For this work, a simple Markov model is used:

$$\begin{bmatrix} \mathbf{r}_{k+1} \\ \mathbf{v}_{k+1} \end{bmatrix} = \begin{bmatrix} 1 & \Delta t \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \mathbf{r}_k \\ \mathbf{v}_k \end{bmatrix} + \begin{bmatrix} \Delta t^2 \\ \Delta t \end{bmatrix} \mathbf{a}_k \tag{12}$$

where $\mathbf{r}_k, \mathbf{v}_k$ are the tracked states at time k, Δt is the integration length and $\mathbf{a}_k \sim \mathcal{N}(\mu_{\mathbf{a}}, \Sigma_a)$ is Gaussian noise that depends on the previous state. Besides these continuous changes, the model accounted for discrete events by stopping the vehicle with some probability if the velocity dropped below a threshold. This simple model was used primarily to demonstrate the use of the tracking algorithms, though more complicated models can be easily implemented (as long as the models are known). Moreover, it is assumed that the parameters of the models are known a priori. In future work, these parameters may be able to be learned/refined through observed data.

3. THEORETICAL RESULTS

3.1 Prior density derivations

The details on the derivations of the prior densities on $\mathbf{p}$ are extensive and are not given here, though a very brief overview is provided. Equations (4) and (5) allow us to deterministically determine two of the six states, conditioned on the other four. For tracking, we have assumed that we know four of these states $(\mathbf{r}_k, \mathbf{v}_k)$ and no approximations are needed. On the other hand, for a detection algorithm, all six stochastic states $(\mathbf{r}, \mathbf{v}, \mathbf{a}$ in two dimensions) must be accounted for. It was found that the prior density conditioned on $\mathbf{r}_0$ is just a linear shifted version when conditioned on $\mathbf{r}_0^*$ near $\mathbf{r}_0$. For the small patches that are considered here (circa 10 m^2 in area), this approximation compared favorably with Monte Carlo simulations. Similarly, by making a far-field approximation, one can linearize the equations into a form that allows for analytical integration. Using these two approximations, one can efficiently compute an unconditioned prior for $\mathbf{p}$ as desired.

3.2 Detection algorithm

For both detection and tracking, it is assumed that we have access to the full phase-history data from which images can be reconstructed in ideal fashion, and that obey the model in Equation (7). For detection, we consider a single road patch at a time with known parameters for describing the random variables $(\mathbf{r_0}, \mathbf{v}, \mathbf{a})$. Thus, it is possible to calculate the prior density on a target signature within that patch, $q(\mathbf{p})$ at any pixel $\mathbf{p}$ in the field of view. Using this information, measurements are made of the form:

$$Z = \sum_{\mathbf{p}} f(\mathbf{p})q(\mathbf{p}) \tag{13}$$

Letting $P_q = \{\mathbf{p} : q(\mathbf{p}) > \epsilon\}$ be the set of pixels with some positive probability, the measurement reduces to the form:

$$Z = \sum_{\mathbf{p} \in P_q} f(\mathbf{p})q(\mathbf{p}) = I_{targ}Q + V \tag{14}$$

where $Q = \sum_{\mathbf{p} \in P_q} q(\mathbf{p})$ is approximately Gaussian with means/variance that can be estimated through Monte Carlo simulation, and $V = \sum_{\mathbf{p} \in P_q} q(\mathbf{p})v(\mathbf{p}) \sim \mathcal{N}\left(0, |P_q|^2\sigma_v^2\right)$. Under this approximation, one can form the likelihood ratio under the hypotheses, $I_{targ} = 0, 1$ as

$$\Lambda(Z) = \frac{\phi(Z, \mu_Q, \sigma_Q^2 + |P_q|^2\sigma_v^2)}{\phi(Z, 0, |P_q|^2\sigma_v^2)} \tag{15}$$

where $\phi(Z; \mu, \sigma^2)$ is the Gaussian density with mean μ and variance σ^2. The test statistic $\Lambda(Z)$ is just the ratio of two Gaussians with different means and variances. Detection performance is evaluated through receiver operating characteristic (ROC) curves based on the values of $\Lambda(Z)$.

3.3 Tracking algorithm

The tracking algorithms take advantage of the prior information in a similar way to the detection algorithm. It should be noted, however, that when we condition on the current state $(\mathbf{r}_k, \mathbf{v}_k)$, the conditional prior density on $\mathbf{p}$, $w_k(\mathbf{p})$ takes on positive probability in only a very few pixels in the reconstructed image. Thus, measurements in the tracking algorithms take the form

$$Z_k = \sum_{\mathbf{p} \in P_w} f(\mathbf{p}) = I_{targ} \sum_{\mathbf{p} \in P_w} \sum_{\tau \in T} \delta[\mathbf{p} - \mathbf{p}^*(\tau)] + \sum_{\mathbf{p} \in P_w} \nu(\mathbf{p}) \tag{16}$$

$$\approx I_{targ}|T| + V. \tag{17}$$

where $P_w = \{\mathbf{p} : w_k(\mathbf{p}) > \epsilon\}$ and $V \sim \mathcal{N}\left(0, |P_w|^2\sigma_v^2\right)$. For small, enough ϵ, the approximation to the first summation is valid by the definition of the prior (assuming the conditioned state to be the actual state). Two filters are implemented to track targets in this situation. The first is a particle filter based on the JMPD particle
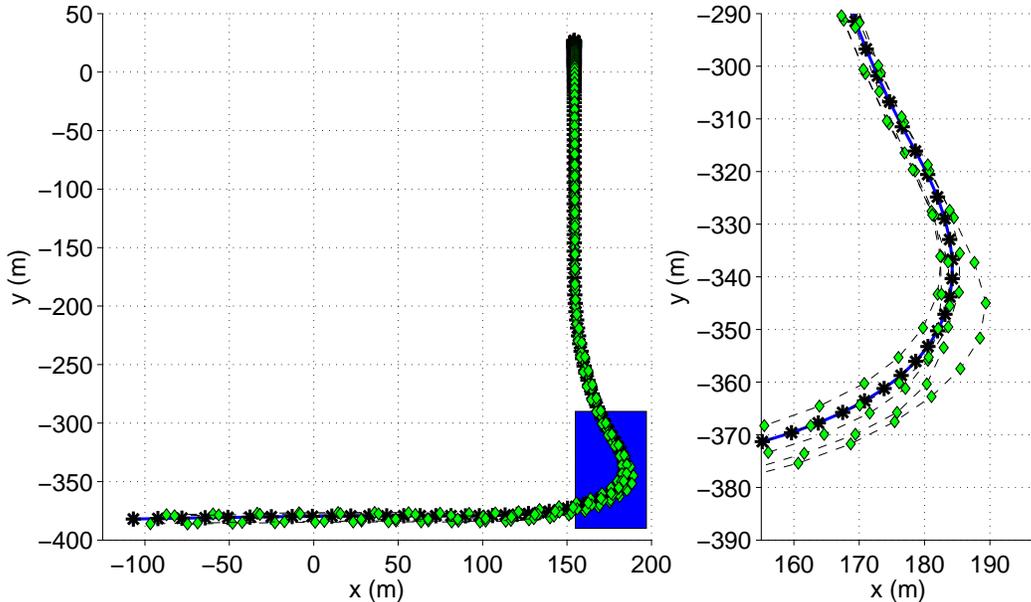
Figure 3. GPS data from the GMTI dataset was used to estimate parameters for a road model. The target's path from GPS data is shown with a dark black line, while green diamonds represent random trajectories of other targets drawn from the fitted road model. The plot on the right zooms in on an interesting region near a bend in the road.

filter.[2] It is well known that the number of particles needed to effectively estimate a state grows very fast with the size of the state. For even a couple of targets, this dimensionality problem would overwhelm the tracking algorithm. The JMPD particle filter addresses this problem by adaptively selecting between two particle filters based on the ambiguity of measurements: when targets are near each other in the sensor space, a coupled particle filter is used; when targets are well separated, an independent particle filter is used to track the targets individually. Assuming that a few states will be coupled at any time, one can approach the speed of using individual particle filters with the accuracy of a fully coupled version.

The unscented Kalman filter (UKF) is also implemented as a baseline comparison. The UKF is an improvement over the standard Kalman filter in cases where the state and/or measurements are highly nonlinear. For the situation addressed in this paper, the state is mostly linear as in Equation (12), except for the case when the target comes to a complete stop (i.e., a discrete event). On the other hand, the measurements are highly non-linear as outlined in Section 2, and the UKF approximates the mean/covariance of the state picks a set of sample points around the mean that captures the true mean/covariance better than the original Kalman filter. Similar to the particle filter, the implemented UKF uses adaptive partitioning of the state based on ambiguity of the measurements in order to efficiently estimate the target states.

4. PERFORMANCE ANALYSIS

To analyze the performance of the aforementioned algorithms, targets trajectories are simulated based on the Gotcha GMTI Data Set[8] which tracked a Durango truck with airborne radar over a period of 71 seconds. Using the Markov model described in Section 2.2, GPS data from the GMTI dataset was used to fit statistical parameters $(\mu_{\mathbf{v}}, \mu_{\mathbf{a}}, \sigma_{vx}^2, \sigma_{vy}^2, \sigma_{ax}^2, \sigma_{ay}^2)$ for individual patches of size 3x3 m^2 along the target's path. Figure 3 shows plots of both the GPS data and random target trajectories drawn from this fitted road model.

4.1 Detection Performance

For both the detection and tracking experiments, between one and five random target trajectories were drawn from among thirty alternatives. For the detection experiment, the ten closest patches to the targets were selected
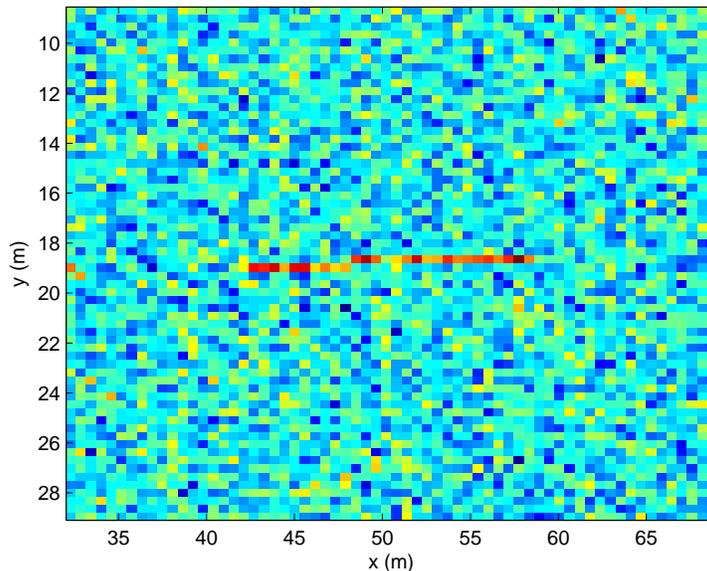
Figure 4. A sample image using ideal reconstruction + Gaussian noise is shown here. Images were 64x64 pixels in size, with each cell containing a minimum area of 0.25m x 0.25m.

at one second intervals. Based on the fitted patch statistics, prior densities were calculated for each patch. SAR images were then created using ideal reconstruction plus Gaussian noise over pixels where the prior density was approximately non-zero. Each image was 64x64 pixels^2 with a minimum area of 0.25m x 0.25m per pixel. A sample image for a single target is shown in Figure 4. Measurements were generated by equation 14 and likelihood ratios were calculated according to equation 15, where μ_Q and σ_Q^2 were estimated by Monte-Carlo simulation. ROC curves were generated from the likelihood ratios over 10 trials, each consisting of 75 time points. The curves are plotted in Figure 5 for various number of random targets. It is seen that for a single target (red solid curve), we have reasonably good detection performance with greater than 0.9 probability of detection for false alarm rates greater than 0.2. However, detection performance decreases when nearby targets are present, which suggests that other processing might be necessary to improve the detection performance.

It is important to note that even if one allowed nearly a 100% false alarm rate, this method still may benefit over other detection algorithms that exhaustively search a field of view. Since we are producing SAR images only in areas where the prior is positive, the computational burden is vastly reduced. Moreover, this algorithm could be used to determine initial guesses for the location of targets that can then be filtered in sequential processing, similar to the tracking algorithms that are discussed next.

A final point with regard to this detection strategy is the implicit geolocation of the target. Other methods for moving target detection using SAR imagery tend to detect targets without any knowledge of the true geolocation of the target. Target geolocation can be provided by the size of the beam or by some beam splitting using multiple phase centers, but the results tend to be coarse with regard to angle. On the other hand, the strategy provided in this paper provides geolocation as a byproduct of detection using only a single phase center. Moreover, this geolocation can be used in fusion with other sensors, potentially enabling sensor management to intelligently track targets.

4.2 Tracking performance

The tracking experiment also selected between one and five random target trajectories. Both filters were updated at 0.1 second intervals, and the particle filter was implemented with 100 particles per target. Measurements at each stage are calculated according to equation 16. Target estimates were initialized to their true values, though in practice initialization may be a key aspect in the performance of these algorithms. Future work plans to analyze robustness to initialization errors and provide possible algorithms to drive down initialization errors.
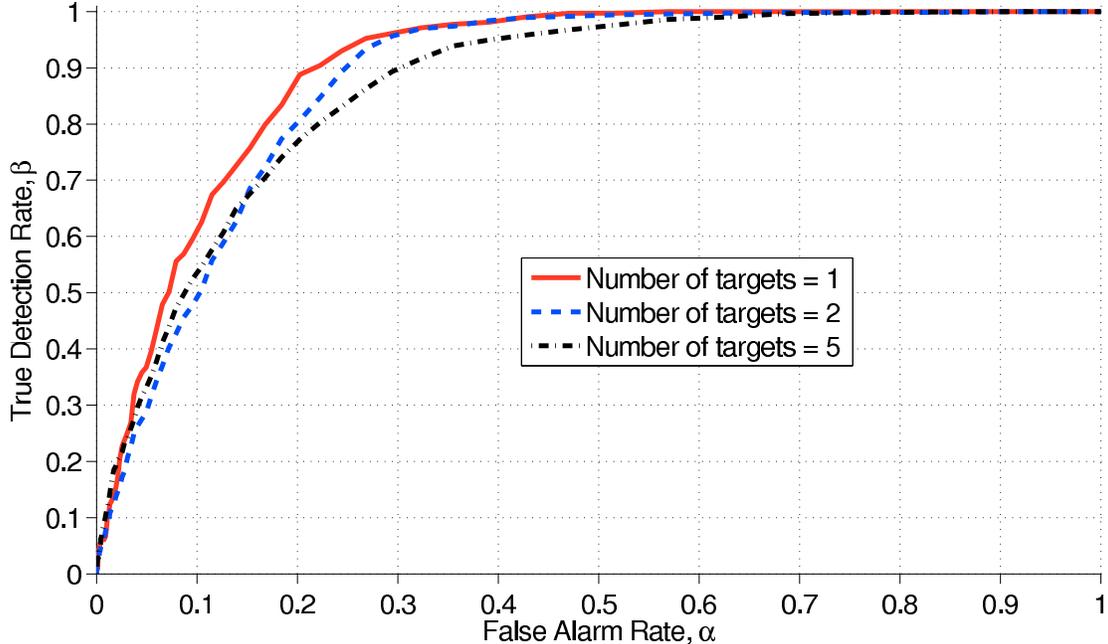
Figure 5. This figure shows the detection performance as ROC curves for various number of targets. It is seen that for a single target (red solid curve), we have reasonably good detection performance with greater than 0.9 probability of detection for false alarm rates greater than 0.2. However, the detection performance decreases as the number of targets grows for most false alarm rates. This suggests that other processing (sequential, track-before-detect, etc.) might be required in order to increase detection performance to an acceptable standard.

	PF		UKF	
Number of Targets	Mean	Std.	Mean	Std.
1	0.41	0.38	3.00	3.59
2	0.47	0.49	5.91	7.36
5	0.63	0.94	5.67	6.44

Table 1. This table shows the mean and standard deviations of the tracking errors in meters for both filters as a function of the number of tracked targets. It is seen that the particle filter (PF) has both lower bias and standard deviations than the unscented Kalman filter (UKF), regardless of the number of targets.

Figure 6 shows the tracking errors from 5 trials for this experiment as function of time. It is clear that the particle filter is more stable than the unscented Kalman filter, which cannot account for discrete events, such as stopping or going off-road. Table 1 shows the mean and standard deviations of the errors for both algorithms as a function of the number of targets. The particle filter has excellent performance that seems relatively consistent as the number of targets grows. This demonstrates that this filter could be used to effectively track a large number of targets. Finally, Table 2 gives the number of targets per partition as a function of the number of tracked targets. It can be seen that the vast majority of states are tracked with nearly independent partitions (62% of states tracked independently for 2 targets, 75% of states tracked with 1 or 2 targets per partition for 5 targets). This suggests that the particle filter can also be efficiently implemented to track multiple targets.

5. CONCLUSIONS AND FUTURE WORK

In this paper, a novel approach to SAR moving target detection was introduced by using prior information on target motion with small regions to enhance target detection/tracking. Prior densities for the target signature within reconstructed SAR images were derived conditioned on various parameters of this target motion model. Moreover, detection and tracking algorithms were provided to take advantage of these priors. It was found that detection performance was reasonable for a single target, but suffered when multiple targets were present.
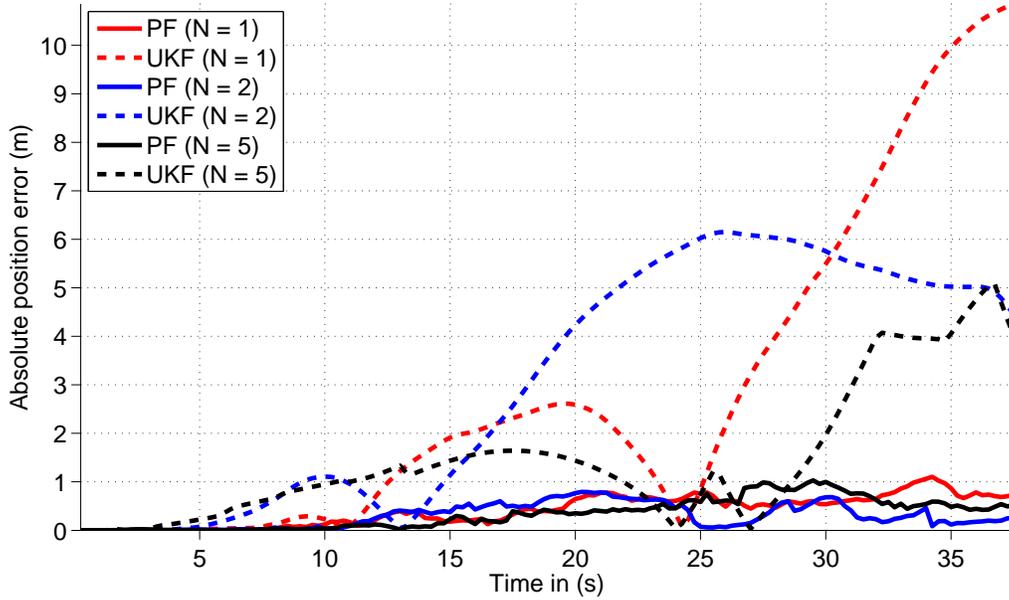
Figure 6. This figure plots the median tracking error (m) for the particle filter (solid lines) and the unscented Kalman filter (dashed lines) at 0.1 second intervals over a period of 38 seconds. Red curves represent experiments with a single target, blue curves represent experiments with two targets, and black curves represent experiments with five targets. The results suggest that the particle filter is superior to the unscented Kalman filter.

	# of targets per partition				
Number of Targets	0	1	2	3	4
1	1.00	0.00	0.00	0.00	0.00
2	0.62	0.38	0.00	0.00	0.00
5	0.44	0.31	0.16	0.08	0.00

Table 2. The number of targets per partition (as a percentage of the number of time points) is shown as a function of the number of tracked targets. It can be seen that the vast majority of stages use nearly independent partitions (either 1 or 2 targets per partition). This suggests that it might be possible to implement an efficient particle filter to track a large number of targets.

Nevertheless, the detection strategy may be advantageous due to its inherent properties that it is computationally efficient and provides geolocation as a by-product of detection. Moreover, track-before-detection algorithms could be implemented to increase detection performance, using provided target tracking algorithm, based on the JMPD particle filter (PF) and the unscented Kalman filter (UKF). It was found that the PF outperformed the UKF in both stability and ability to track multiple targets.

Future work plans to learn both the target motion model within a single road patch, as well as the Hidden Markov Model for updating the target state over time. Moreover, we plan to explore robustness to the PF initialization, and possibly introduce algorithms to provide good initial estimates. Finally, we would like to explore performance bounds for estimating a target state based on parameters we have control over (such as integration time, number of pixels in the reconstructed image, etc.), and subsequently determine an optimal strategy based on this knowledge.

6. ACKNOWLEDGMENTS

REFERENCES

[1] J. Jao, "Theory of synthetic aperture radar imaging of a moving target," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 39, no. 9, pp. 1984–1992, 2001.

[2] C. Kreucher, K. Kastella, and A. O. H. III, "Multitarget tracking using the joint multitarget probability density," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 41, no. 4, pp. 1396–1414, October 2005.

[3] E. Wan and R. Van Der Merwe, "The unscented Kalman filter," *Kalman filtering and neural networks*, pp. 221–280, 2001.

[4] J. Fienup, V. Syst, and M. Ann Arbor, "Detecting moving targets in SAR imagery by focusing," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 37, no. 3, pp. 794–809, 2001.

[5] L. A. Gorham and L. J. Moore, "SAR image formation toolbox for MATLAB," E. G. Zelnio and F. D. Garber, Eds., vol. 7699, no. 1. SPIE, 2010.

[6] S. Werness, W. Carrara, L. Joyce, and D. Franczak, "Moving target imaging algorithm for SAR data," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 26, no. 1, pp. 57–67, 1990.

[7] M. Kirscht, D. GmbH, and G. Friedrichshafen, "Detection and imaging of arbitrarily moving targets with single-channel SAR," *IEE Proceedings-Radar, Sonar and Navigation*, vol. 150, no. 1, pp. 7–11, 2003.

[8] S. M. Scarborough, C. H. Casteel Jr, L. Gorham, M. J. Minardi, U. K. Majumder, M. G. Judge, E. Zelnio, M. Bryant, H. Nichols, and D. Page, "A challenge problem for SAR-based GMTI in urban environments," E. G. Zelnio and F. D. Garber, Eds., vol. 7337, no. 1. SPIE, 2009, p. 73370G. [Online]. Available: http://link.aip.org/link/?PSI/7337/73370G/1