

Switching Kalman Filters for Prediction and Tracking in an Adaptive Meteorological Sensing Network

Victoria Manfredi, Sridhar Mahadevan, Jim Kurose
Department of Computer Science
140 Governors Drive
University of Massachusetts
Amherst, Massachusetts 01003-4601
Email: {vmanfred, mahadeva, kurose}@cs.umass.edu

Abstract— We consider the problem of configuring sensors in an adaptive sensor network being used to monitor meteorological features. One way to decide future sensor configurations is to base them on information currently being collected. For instance, if a meteorological sensor network is being used to monitor storms in Oklahoma, then the sensors could be dynamically configured based on the predicted storm locations. While Kalman filters and their extensions are commonly used for prediction and tracking, they have been primarily applied to objects with known or fixed dynamics such as missiles or people. We explore the advantages and limitations of using Kalman filters to track objects with nonstationary dynamics (e.g., a storm can grow in size). In particular, we focus on tracking meteorological features over time with the objective of using this information to determine where radars should focus their sensing. We present results for tracking storm cells comparing least-squares regression with Kalman filter and switching Kalman filter methods. Our results show that on average the Kalman filter methods better predict the future location of a storm centroid than does a least-squares regression algorithm currently in use for meteorological storm tracking.

I. INTRODUCTION

A sensor network is comprised of sensors collecting data in an environment. An adaptive sensor network is a sensor network in which the sensor configurations can be changed, either by users or under system control. For example, the network might change the frequencies at which measurements are taken, or change the locations where *in situ* or remotely sensed measurements are taken. The ability to dynamically change sensor configurations provides a great deal of flexibility in how, when, where, and what information is collected. The utility of this adaptability is most apparent when the sensor network's

resources are limited, such as when all locations cannot be sensed all of the time.

One approach for determining future sensor configurations is to base them on information currently being collected or that has been collected in the past. For instance, suppose an adaptive radar network is being used to monitor storms in Oklahoma. The future paths of the storms being monitored can be predicted based on current and previous monitored storm data. Then if a storm is predicted at a certain latitude and longitude in five minutes, the radars can be configured to scan that location at that time.

In this paper we examine the problem of tracking storm cells over time with the objective of using current and past observations to intelligently determine where radars should focus their sensing. The context for this research is the CASA project (Collaborative Adaptive Sensing of the Atmosphere) [31]. Although the CASA distributed radar network will monitor a variety of meteorological features, including tornadoes, storm cells and shear, this paper focuses only on tracking storm cells. We note that our focus on tracking only storm cells is due to a lack of track data for tornadoes and other meteorological features. Tornado tracking and prediction would be very interesting and we hope that this work can be extended to such problems.

This paper examines the use of the Kalman filter and the switching Kalman filter for tracking storm cells. There are several advantages to using Kalman filters. First, prediction using the basic Kalman filter is extremely fast and requires little memory. This is essential for the real-time requirements of the CASA system. Second, an error estimate is associated with each prediction. Third, these predictions can be computed

recursively, bounding the time and memory needed for computation. While the Kalman filter and its extensions are commonly used for prediction and tracking, they have been primarily applied to objects with known or fixed dynamics such as missiles or people. In order to use Kalman filters to track storms (or other meteorological features), several potential problems must be addressed. First, storms can change in size and intensity; they can also split into separate storms, or merge with other storms. Second, different storms do not necessarily have the same dynamics and an individual storm’s dynamics will likely not satisfy the linear-Gaussian assumption of the Kalman filter. Thus, one of the goals of this paper is to examine the performance of the Kalman filter for tracking meteorological features. Our results show that on average the Kalman filter methods better predict the future location of a storm centroid than does a least-squares regression algorithm currently in use for meteorological storm tracking.

The rest of this paper is organized as follows. Section 2 gives background on the prediction methods we use. Section 3 reviews current tracking methods. Section 4 formulates the problem we are addressing and describes our Kalman filter-based approach. Section 5 presents experimental results for storm cell tracking comparing least-squares regression, Kalman filter, and switching Kalman filter methods. Finally, Section 6 summarizes important issues identified in this project and Section 7 outlines future work.

II. BACKGROUND

This section provides a brief review of the graphical models framework, and explains the basics of Kalman filters, switching Kalman filters, and inference.

A. Graphical Models

See [20], [27] for a more comprehensive treatment of graphical models. A graphical model efficiently encodes a joint probability distribution using a graph. Nodes in the graph represent random variables; edges in the graph capture dependencies between variables. Based on the dependencies being encoded, edges may be directed, as in a Bayesian network (BN), or undirected as in a Markov field. Each node that represents a discrete random variable has an associated conditional probability table (CPT). Each node that represents a continuous random variable has an associated conditional probability distribution (CPD). For all possible node and parent values, the CPT or CPD for a given node contains the




Fig. 1. DBN formulation of the Kalman filter. X represents state and Y represents observations. Clear nodes are hidden and shaded nodes are observed.

probability that the node has some value conditioned on a set of values for its parents.

The Kalman filter and switching Kalman filter can be represented by models in one class of directed graphical models, specifically dynamic Bayesian networks (DBN). In a DBN, the values of the variables change over time according to a (probabilistic) transition model; while the values of the variables change, the transition model does not. A DBN can be thought of as a BN repeated for multiple time steps, with additional edges to connect variables at time t with variables at time $t + 1$. In both BNs and DBNs, the problem of inference is to determine the probability distribution for a query variable (or variables) given some evidence. For instance, a DBN with variables x and y could be queried for the probability that variable x takes on some value at time t given observed values of y for time steps 1 to t .

B. Kalman Filter

See [11], [20], [27] for a more comprehensive treatment of Kalman filters. Kalman filters (KF) have traditionally been used in tracking. Unlike in least-squares regression, the KF maintains a state representation that is updated on each timestep. This makes the KF a type of state-space model. In a KF, the true location \mathbf{x}_t of the object is assumed hidden and is modeled as a Gaussian random variable. Noisy observations \mathbf{y}_t of the hidden state are assumed available. Transitions from one hidden state to another are modeled with a linear Gaussian function: the next location \mathbf{x}_{t+1} is a linear function of the current location \mathbf{x}_t plus some Gaussian noise \mathbf{w}_t where $\mathbf{w} \sim N[0, Q]$. Observations are also assumed to be a linear function of the hidden state plus some Gaussian noise \mathbf{v}_t where $\mathbf{v} \sim N[0, R]$ and $\mathbf{y}_1 \sim N[\mu, \Sigma]$. That is,

$$\begin{aligned} \mathbf{x}_{t+1} &= A\mathbf{x}_t + \mathbf{w}_t \\ \mathbf{y}_t &= B\mathbf{x}_t + \mathbf{v}_t \end{aligned}$$

Using the observations \mathbf{y}_t , the KF computes the probability of state \mathbf{x}_t given the observations up to time t , that is $P(\mathbf{x}_t | \mathbf{y}_1, \dots, \mathbf{y}_t)$. This is called filtering. The

probability of a future state \mathbf{x}_{t+k} given the current state \mathbf{x}_t , that is $P(\mathbf{x}_{t+k} | \mathbf{x}_t)$ can also be computed. This is called prediction. Finally, given future observations up to time $t + k$, the probability of the state at time t , $P(\mathbf{x}_t | \mathbf{y}_1, \dots, \mathbf{y}_{t+k})$, can be computed. This is called smoothing. Filtering, prediction, and smoothing are all types of inference.

One of the major advantages of the KF is that the number of parameters needed to represent the state space does not explode as more observations are acquired. This is because the KF models the state as a Gaussian random variable and state transitions as a linear Gaussian function. Consequently, the state remains a Gaussian random variable, requiring only that the mean and the covariance be stored. This permits the KF to be compactly formulated using recursive updates. However, the dynamics of the object being tracked may not exactly satisfy the linear Gaussian assumption. For instance, the noise in the observations may not be Gaussian, or the object may be moving in a nonlinear fashion. The extended Kalman filter, the unscented Kalman filter, and the switching Kalman filter address the assumption of linear dynamics while particle filters additionally address the assumption of Gaussian noise. Figure 1 shows the DBN representation for the KF.

C. Switching Kalman Filter

See [2], [6], [12], [19] for a more comprehensive treatment of switching Kalman filters. Switching Kalman filters (SKF), also known as jump Markov linear systems, maintain a state representation that is updated on each timestep. Like with the KF, this also makes the SKF a type of state-space model. The basic difference between the SKF and KF is that the SKF can track objects with nonlinear dynamics while the KF can only track objects with linear dynamics. For instance, suppose that a KF is being used to track an object and the object is moving according to a dynamical model that can be approximated by the KF. Suppose that something then causes the dynamics of the object’s movement to change, such as having to avoid an obstacle. A KF will initially predict poorly since the dynamics of the object have changed and a single KF can only encode one dynamical model. If it were possible, however, to recognize (or predict) when the object’s dynamics were about to change, a different KF (encoding the new dynamics) could be used to track the object at that point. The SKF attempts to do exactly that. Essentially, an SKF computes a piece-wise linear approximation of the path of an object moving with nonlinear dynamics. In an SKF,




Fig. 2. DBN formulation of the switching Kalman filter. S represents switches, X represents state, and Y represents observations. Clear nodes are hidden and shaded nodes are observed.

the value of a switch variable indicates the particular KF being used on each time step. Like KFs, SKFs can also be formulated as a DBN. Figure 2 shows the DBN representation for the SKF.

In an SKF, both the values of the state variable X and the values of the switch variable S are unknown. Thus, the KF that is generating observations at a given timestep is unknown. As a consequence, unlike in the KF, the problem of inferring the switch node or hidden state values in a SKF is intractable. For example, suppose we are observing an object whose dynamics can be described by switching among n KFs. At timestep 1, it is unknown which KF generated the observation, so the state distribution for each KF must be maintained, giving a belief state of size n . At timestep 2, which KF generated the observation is again unknown. Because the KF for timestep 1 is also unknown, the belief state is now of size n^2 . Consequently, at timestep T , the belief state will be of size n^T . Thus, inference based on an exact state description is intractable for SKFs. Various approximate inference techniques have, however, been developed to address the problem of inference in SKFs and in general. We will focus on a method called the Generalized Pseudo Bayesian Algorithm [2], [12].

D. Generalized Pseudo Bayesian Algorithm

See [2], [12] for a more comprehensive treatment of the Generalized Pseudo Bayesian algorithm. The Generalized Pseudo Bayesian algorithm is an approximate inference method for the SKF. The essential idea of many of the approximate inference methods for SKFs is that if the values of the switch variables in the SKF were known, then the belief state would no longer be exponential in size and the Kalman filter equations could be used to efficiently compute the values of the state variables.

In comparison, rather than explicitly estimating switch

node values, the Generalized Pseudo Bayesian algorithm works by weighting and summing over the hidden state values of each Kalman filter for those values that are more than k timesteps old. The weights are the probabilities that each Kalman filter generated the observation k timesteps ago. The number of timesteps k can be chosen, leading to an order k Generalized Pseudo Bayesian algorithm. Of course, the larger the k , the greater the cost of inference. Suppose $k = 2$ and let \mathbf{x}^{ij} be the belief state if Kalman filter i is used on timestep $t - 1$ and Kalman filter j is used on timestep t , let V^{ij} be the corresponding covariance, and let W^{ij} be the probability that the observation at timestep $t - 1$ was generated using Kalman filter i and the observation at timestep t was generated using Kalman filter j . Then the order two generalized pseudo-Bayesian algorithm works by “collapsing” over the means and covariances two time steps ago as follows.

$$\begin{aligned} (\mathbf{x}^j, V^j) &= \text{Collapse}(\mathbf{x}^{ij}, V^{ij}, W^{ij}) \\ \mathbf{x}^j &= \sum_i W^{ij} \mathbf{x}^{ij} \\ V^j &= \sum_i W^{ij} V^{ij} + \sum_i W^{ij} (\mathbf{x}^{ij} - \mathbf{x}^j)(\mathbf{x}^{ij} - \mathbf{x}^j)' \end{aligned}$$

Other approximate inference methods for the SKF besides the Generalized Pseudo Bayesian algorithm include variational inference, sampling methods, and a Viterbi algorithm [20], [24].

III. RELATED WORK ON TRACKING

We distinguish general statistical tracking techniques from storm tracking algorithms which may have a statistical foundation. Storm tracking methods can be divided into extrapolation and/or statistical algorithms, and other more knowledge-intensive algorithms. Extrapolation algorithms use deterministic linear models and can be loosely divided into centroid methods and correlation methods. Centroid methods work by identifying and tracking features associated with storm cells. The extrapolation algorithms have low complexity; their disadvantages are that they track the storm cell rather than the larger storm, resulting in less accurate long-term predictions, and that they have difficulty identifying when and how storms split.

The SCIT Algorithm [10] is one example of an extrapolation algorithm. It uses linear-least squares over the last five data points to predict storm locations; data points are the centroid locations of the storm cell being tracked. The TITAN methodology of [4] is another extrapolation algorithm, but it also uses cross-correlation.

More knowledge-intensive algorithms for storm tracking include the Gandolf system developed by [25] which treats storm cells as objects and then models the meteorologic evolution of each object; and the Growth and Decay Storm Tracker [29] which focuses on tracking the encompassing storm using an elliptical filter, rather than tracking a storm cell.

Unlike the approaches discussed above, rather than working with storm centroids, [13], [14] work with radar images. Given a sequence of weather images (e.g., radar reflectivity), [13], [14] first use K-means to identify storm clusters at different levels of granularity. The motion estimates for clusters in an image at time t are then computed by finding the corresponding regions of pixels for the image at time $t + 1$ that minimize error. These estimates are then smoothed using a Kalman filter. Computing these estimates at different levels of granularity has the advantage of permitting both longer-term and shorter-term forecasts.

In comparison, the Ensemble Kalman filter [5], [9], [18], [22] maintains a state representation of the atmosphere and evolves the state according to atmospheric dynamics. A single state consists of (*latitude, longitude*) and the accompanying features for that location such as temperature and surface pressure for different atmospheric levels. Because the state represents the atmosphere, the size of the state vector is large, and estimating the covariance using the standard Kalman filter is computationally expensive. To address this, ensemble-based techniques like the Ensemble Kalman filter were developed in which an ensemble of trajectories are used to estimate the covariance. Unlike the Kalman filter, the Ensemble Kalman filter can approximate nonlinear dynamics. Increasing the number of ensemble members in the Ensemble Kalman filter decreases the root mean square error in the forecast but increases the computational cost [9], [22].

Work on tracking within the machine learning community is primarily statistical: techniques such as Kalman filters, extended Kalman filters, switching Kalman filters, and particle filters in particular, are commonly used. Traditional tracking applications include tracking of ballistic objects such as missiles, and tracking of enemy aircraft [26]. Other applications include tracking vehicles [15] and tracking people [21], [24]. For instance, switching Kalman filters have been used to track humans alternating between running and walking [24].

While tracking in sensor networks is based on statistical techniques, like the Kalman filter [21] or the extended Kalman filter [3], the distributed aspect of the

sensor network introduces interesting variations on the tracking problem such as distributed versus centralized tracking [3]. In sensor networks comprised of many low-power sensors, work on tracking focuses on how to leverage the distributed aspects of the network to, for instance, minimize power consumption or distribute computation. In [16], a dynamically changing subset of the sensors (a “collaborative group”) is maintained for each object being tracked, but only the current group leader performs tracking, thereby minimizing energy costs. Similarly, [23] uses collaborative groups of video sensors to track targets such as vehicles. This contrasts with the CASA [31] radar network: the key idea for using a distributed collaborative adaptive radar network in CASA is that the more radars that scan the same region, the more informative will be the collective data. Hence, the current goal of collaboration in the CASA network is not to decrease energy costs but rather to improve detection of sensed phenomena.

The adaptivity of the CASA radar network, however, brings it back into the realm of problems encountered in sensor networks. Various meteorological features of interest and differing end user interests in the various features in the environment give rise to a resource allocation problem in which the targeting of the radars is to be optimized. Despite this difference from the case of power-constrained sensor networks, collaborative groups as in [8] or [16] could potentially be formed to solve smaller local optimization problems.

IV. COMPARISON OF KALMAN FILTER, SWITCHING KALMAN FILTER, AND LEAST-SQUARES REGRESSION

This section presents results comparing how well the KF, the SKF, and linear least-squares regression (SCIT [10]) perform in tracking a storm and predicting its future location. We constrain the problem in two ways. First, we are only interested in tracking a previously identified storm cell; we assume that there are meteorological algorithms that perform identification. Second, we assume that storm cell observations are obtained at regular intervals. This is not a necessary assumption, but it simplifies prediction using the Kalman filter.

A. Implementation

We implemented the SKF equations found in [19] in C++ and using the matrix library CwMtx and the GNU Scientific library. The SKF model was used to do prediction in both the basic KF (by setting the number of

KFs to one) and the SKF. The regress function in Matlab was used for linear least-squares regression.

B. Data

Real storm track data was obtained from NSSL (National Severe Storms Laboratory), courtesy of Kurt Hondl and the WDSS-II software [7]. The data consists of 35 storm tracks ranging in length from ten to 30 data points, identified using the SCIT algorithm [10]. Each track is a sequence of latitude and longitude coordinates with observations occurring approximately every five minutes. We note that SCIT, which stands for Storm Cell Identification and Tracking, both identifies and tracks storms. Each new datapoint that the SCIT algorithm acquires must be assigned to either an existing storm track or to a new storm track. A regression model is then fit to the five most recent datapoints in the track and used to predict future datapoints for that track. Our data consists of the storm tracks produced by SCIT.

C. Inference and Prediction

Inference in the SKF was done using a second order Generalized Pseudo Bayesian algorithm: i.e., collapsing was done over states two timesteps and older. This required computing the terms $\Pr(S_{t-1} = i, S_t = j|y_{1:t})$ and $\Pr(S_t = j|y_{1:t})$, where S_t represents the switch node value at time t and $y_{1:t}$ represents the sequence of observations from time 1 to t . These terms permit us to compute the most likely sequence of switch node values with $\operatorname{argmax}_j P(S_t = j|S_{t-1} = i, e_{1:t})$ beginning at time 1 and ending at time t . In comparison to the Viterbi algorithm (see [20], [27] for a description), this algorithm is simpler because we have already computed the terms $\Pr(S_{t-1} = i, S_t = j|y_{1:t})$ and $\Pr(S_t = j|y_{1:t})$. The switch node value at t indicates the KF that will give the best prediction, and hence that KF is used to do the 1-step prediction. Prediction with linear least-squares was done by using the previous five points to compute the line with the smallest least-squares error. The average x -interval was then added to the most recent x -coordinate to get the next x -coordinate. Least-squares was then used to determine the corresponding y -coordinate.

D. Training

We compared hand-initializing the parameters with using Expectation-Maximization (EM) [19] to learn the parameters for the KF and the SKF models. See [19] for an explanation of the model parameters. We defined each observation to be a vector of length two comprised of a (*latitude, longitude*) pair. We defined the mean of

each hidden state to be a vector of length four comprised of the 4-tuple, (*latitude, longitude, latitude velocity, longitude velocity*). We hand-coded the parameters as follows. The KF used generated points to the northeast. That is, the parameter matrix A was,

$$A = \begin{bmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

The first column represents latitude, the second represents longitude, the third represents latitude velocity, and the fourth represents longitude velocity. The first row encodes the linear update equation for the latitude, the second row encodes the update equation for the longitude, the third row encodes the update equation for latitude velocity, and the fourth row encodes the update equation for longitude velocity. For instance, the latitude at time $t + 1$ will be the latitude for the hidden state at time t , since there is a 1 in row 1 column 1, plus the latitude velocity for the hidden state at time t , since there is a 1 in row 1 column 3. Similarly, longitude at time $t + 1$ will be the longitude for the hidden state at time t , since there is a 1 in row 2 column 2, plus the longitude velocity for the hidden state at time t , since there is a 1 in row 2 column 4. Consequently, if there is little noise and no observations, then iterating this KF will generate points to the northeast.

The SKF used four KFs; three generated points to the north, east, and northeast respectively and the fourth encoded a stationary system. For both the KF and the SKF, the first observation y_1 was used for the initial state μ . For the SKF, the prior and transitions for the switches were uniform, i.e., all KFs had equal probability of being used. Below are the parameters for the SKF. The KF parameters are those for $s = 1$. The remaining hand-coded parameters for all $1 \leq s \leq 4$ are,

$$\begin{aligned} \Sigma_s &= \begin{bmatrix} 10 & 0 & 0 & 0 \\ 0 & 10 & 0 & 0 \\ 0 & 0 & 10 & 0 \\ 0 & 0 & 0 & 10 \end{bmatrix} & Q_s &= \begin{bmatrix} 0.1 & 0 & 0 & 0 \\ 0 & 0.1 & 0 & 0 \\ 0 & 0 & 0.1 & 0 \\ 0 & 0 & 0 & 0.1 \end{bmatrix} \\ R_s &= \begin{bmatrix} 0.01 & 0 \\ 0 & 0.01 \end{bmatrix} & B_s &= \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \\ A_1 &= \begin{bmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} & A_2 &= \begin{bmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \\ A_3 &= \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} & A_4 &= \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \end{aligned}$$

We now present the learned KF parameters obtained using EM. Due to the limited number of tracks, parameter learning was done with leave-one-out cross-validation: i.e., we trained on 34 tracks and tested on the 35th. We did this 35 times so that we could test on each of the tracks. The learned KF parameters below are from a representative training run on 34 tracks.

$$\begin{aligned} \mu &= \begin{bmatrix} 32.2006 \\ -94.58 \\ 0.06057 \\ -0.00713 \end{bmatrix} & R &= \begin{bmatrix} 0.00885 & 0.00124 \\ 0.00124 & 0.00663 \end{bmatrix} \\ B &= \begin{bmatrix} 0.9987 & -0.0005 & -0.0068 & -0.0068 \\ -3.1034e-06 & 1.0001 & -0.0131 & 0.0130 \end{bmatrix} \\ \Sigma &= \begin{bmatrix} 10.5616 & -6.4091 & -0.0106 & 0.0478 \\ -6.4091 & 4.4895 & 0.0134 & -0.0338 \\ -0.0106 & 0.0134 & 9.0311e-05 & -0.0001 \\ 0.0478 & -0.0338 & -0.0001 & 0.0003 \end{bmatrix} \\ Q &= \begin{bmatrix} 0.04520 & 0.01209 & -0.00075 & -0.00074 \\ 0.01209 & 0.02358 & -0.00279 & 0.00204 \\ -0.00075 & -0.00279 & 0.07144 & 0.00319 \\ -0.00074 & 0.00204 & 0.00318 & 0.06414 \end{bmatrix} \\ A &= \begin{bmatrix} 0.99518 & -0.00177 & 0.30253 & 0.05275 \\ 0.00311 & 1.00081 & 0.05589 & 0.20061 \\ -0.00354 & -0.00121 & 0.47183 & -0.01887 \\ 0.001878 & 0.00049 & 0.00159 & 0.48156 \end{bmatrix} \end{aligned}$$

The μ and Σ parameters represent the mean and covariance of the initial state respectively; the first row/column represents latitude, the second represents longitude, the third represents latitude velocity, and the fourth represents longitude velocity. The Q parameters represent the covariance for hidden state while the R parameters represent the covariance for the observations.

E. Results

Table I summarizes the average 1-step and 2-step root mean-square error (RMSE) for the latitude and longitude error over all storm tracks for each of the methods. Table II contains the individual latitude and longitude 1-step prediction RMSEs for each storm track for each of the methods. The RMSE was calculated using only predicted points. For 1-step prediction, the first predicted point was the 6th point; for 2-step prediction, the first predicted point was the 7th point. Figure 3 plots the 1-step predictions for linear least-squares regression and the hand-coded SKF for three different storm tracks (see [17] for more storm track plots). Table III summarizes the average computation time required for making a 1-step prediction. Computation times were recorded for a 1 GHz Dell Optiplex GX150.

TABLE I

SUMMARY OF THE PREDICTION ERROR. EM INDICATES THAT THE MODEL PARAMETERS WERE LEARNED.

Method	Avg 1-Step RMSE		Avg 2-Step RMSE	
	Latitude	Longitude	Latitude	Longitude
KF	0.2248	0.1923	0.3359	0.2871
KF-EM	0.1967	0.1680	0.2680	0.2389
SKF	0.1914	0.1702	0.2642	0.2421
SKF-EM	0.2577	0.2070	0.3948	0.3110
LS	0.2114	0.2107	0.3030	0.3081

TABLE II

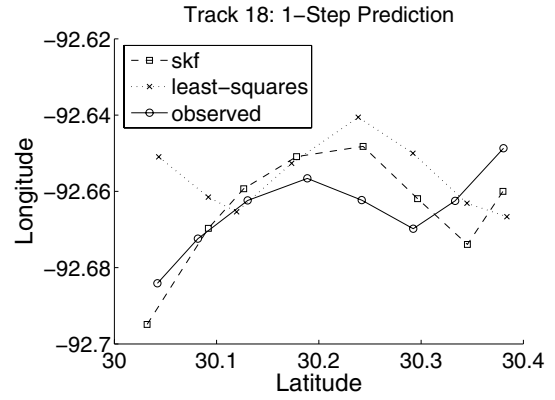
ROOT MEAN SQUARE ERROR OF INDIVIDUAL TRACKS FOR 1-STEP PREDICTIONS. COLUMNS ARE LATITUDE LONGITUDE.

Track (length)	KF-EM		SKF		Least-Squares	
1(10)	0.0669	0.0286	0.0288	0.0285	0.0308	0.0524
2(10)	0.5668	0.5763	0.6031	0.5938	0.5381	0.5516
3(10)	0.4658	0.4596	0.5486	0.4488	0.5886	0.4599
4(10)	0.4693	0.0532	0.4837	0.0677	0.4879	0.0916
5(10)	0.0740	0.0191	0.0211	0.0281	0.0191	0.0201
6(10)	0.0947	0.0200	0.0230	0.0270	0.0202	0.0203
7(10)	0.0321	0.0197	0.0173	0.0010	0.0190	0.0080
8(10)	0.0347	0.0387	0.0298	0.0209	0.0302	0.0137
9(10)	0.6926	0.6271	0.7653	0.5751	0.7024	0.7652
10(10)	0.0610	0.5968	0.0683	0.5821	0.0666	0.5813
11(12)	0.0297	0.0244	0.0197	0.0279	0.0162	0.0388
12(12)	0.0135	0.0087	0.0251	0.0105	0.0251	0.0111
13(12)	0.0830	0.4339	0.0668	0.4413	0.0631	0.4579
14(13)	0.9485	0.0241	0.9563	0.0350	0.9575	0.0335
15(13)	0.0256	0.0574	0.0326	0.0585	0.1574	0.2630
16(13)	0.3543	0.4235	0.3606	0.4143	0.3876	0.4645
17(13)	0.1490	0.5418	0.1563	0.5547	0.1577	0.9725
18(13)	0.0271	0.0350	0.0079	0.0092	0.0088	0.0174
19(13)	0.0148	0.0135	0.0118	0.0092	0.0135	0.0674
20(13)	0.0144	0.0221	0.0151	0.0165	0.0134	0.0157
21(14)	0.1546	0.0465	0.1477	0.0448	0.1457	0.0550
22(14)	0.0584	0.0393	0.0268	0.0282	0.0227	0.0623
23(14)	0.0243	0.0125	0.0157	0.0110	0.0141	0.0146
24(15)	0.0654	0.0272	0.0391	0.0281	0.0362	0.0514
25(15)	0.1394	0.0283	0.1307	0.0186	0.1304	0.0190
26(15)	0.0192	0.0143	0.0171	0.0158	0.0532	0.1206
27(16)	0.0327	0.1051	0.0333	0.1088	0.0352	0.1140
28(16)	0.3871	0.2127	0.3918	0.2128	0.4226	0.2576
29(18)	0.1505	0.2664	0.1581	0.2683	0.1614	0.2859
30(19)	0.7902	0.3782	0.8126	0.3854	0.8863	0.4001
31(21)	0.0651	0.0332	0.0158	0.0463	0.0139	0.0325
32(25)	0.0361	0.0329	0.0129	0.0236	0.0129	0.0298
33(26)	0.2318	0.1773	0.2400	0.1785	0.2625	0.1943
34(29)	0.6633	0.6299	0.6830	0.6499	0.8844	0.8105
35(30)	0.0457	0.0215	0.0142	0.0213	0.0132	0.0200


TABLE III

TIME REQUIRED TO MAKE A 1-STEP PREDICTION ACCORDING TO THE NUMBER OF KALMAN FILTERS USED.


# of KFs	Time for 1-Step Prediction in Seconds		
	Average	Maximum	Minimum
1	0.000155	0.000864	0.000103
4	0.001689	0.006479	0.001138
8	0.006655	0.028375	0.004553



(a) Track 18



(b) Track 31



(c) Track 33

Fig. 3. 1-Step predictions for storm tracks 18, 31, and 33. Switching Kalman filter parameters were hand-coded.

F. Discussion

Although linear least-squares regression (SCIT [10]) was originally used to identify the storm tracks from raw data, Table I shows that the average track RMSE is lowest for the learned KF and hand-coded SKF models. Similar results hold for both the 1-step and 2-step prediction. Because of the limited amount of training data, we will focus on the results for the hand-coded SKF. In comparison to Table I, Table II shows that the individual track RMSE is sometimes lower for the SKF and sometimes lower for the least-squares method. Additionally, the SKF may have lower latitude error but higher longitude error, or vice versa. We therefore compare the different methods based on the track RMSE instead of the average RMSE. For instance, Figure 3a shows 1-step predictions for track 18 for which the SKF has lower RMSE, while Figure 3b shows 1-step predictions for track 31 for which least-squares has lower RMSE. The most apparent difference is that the track in Figure 3a is much less linear than the track in Figure 3b, which would explain why least-squares had higher RMSE than the SKF for track 18.

Examining the prediction errors in Table II shows that for a given method, the latitude RMSE is often higher than the longitude RMSE, particularly for the Kalman filter methods. This may be a function of the storm dynamics or measurement technique, since the discrepancy is most apparent for tracks with large RMSE such as track 33 (Figure 3c). We note that the least-squares method seems to have less of a bias between the latitude and longitude RMSEs. This is particularly apparent if the average RMSEs for least-squares in Table I are examined. It should be noted that 0.1 degrees of latitude corresponds to 6.9 miles and that 0.1 degrees of longitude corresponds to at most 6.9 miles [1].

The distribution of the prediction errors in Table II for the SKF and least-squares varies widely, from about 0.88 to about 0.008. Examining the actual tracks shows that tracks with large errors, such as track 33 in Figure 3c, often looked like they were two different tracks joined together. The tracks were originally identified using SCIT so this phenomenon may be a result of how new datapoints were associated with existing tracks. The SKF framework could be used instead of SCIT to associate new datapoints with existing tracks or to identify new tracks.

The KF using learned parameters does well, about as well as the SKF using hand-coded parameters. This is one indication that the poor performance of the KF

using hand-coded parameters is due in part to improperly initialized parameters. One caveat is that there were a limited number of tracks for training, hence the use of leave-one-out cross-validation. More tracks would help verify the performance of the KF and SKF with learned parameters. We note that the SKF using learned parameters performs worse than the KF with hand-coded parameters. This may be due in part to overfitting as a result of the small dataset. Work by [6] discusses the difficulties with training SKFs (referred to in that paper as switching state-space models), in particular the difficulty in learning when to switch. We found empirically that the performance of the learned models depended strongly on how the model parameters were initialized before learning. For instance, although not shown, we found that with a slightly different parameter initialization (using $0.01 \cdot I$ instead of $0.1 \cdot I$ for the Q matrices where I is the identity matrix), and using 8 KFs in the SKF, gave improved performance for the learned SKF over the hand-coded SKF. Also note that the Kalman filter models need to condition on an adequate number of observations occurring sufficiently frequently to do well. With only a few observations or with infrequent observations, linear least-squares regression may have the advantage. We note that linear least-squares regression over only the last five data points is a sort of piece-wise linear approximation. But unlike in the SKF model, linear least-squares regression can only switch to an arbitrarily sloped new line segment every five points.

Finally, Table III shows that while increasing the number of Kalman filters used in the SKF increases the computation time required for making a 1-step prediction, the computation time required is still sub-second. Although not shown, we also found that (as expected) the computation time required for making a 1-step prediction is not affected by the number of previous observations. With respect to other storm-tracking methods besides least-squares, we note that the hierarchical clustering method in [13] uses data acquired every 5-6 minutes and requires thirty seconds of computation to perform prediction. In comparison, the ensemble Kalman filter makes long-term (12 hour) forecasts [18].

V. SUMMARY

For certain types of tracks, the SKF appears to have the potential to better predict the future location of a storm centroid than the linear least-squares SCIT [10] algorithm. Since we compared the SKF and linear least-squares regression on data already processed by SCIT, a comparison of the two methods on raw data would

be useful. Some of the difficulties initially anticipated with tracking storms were not encountered. For instance, storms are not discrete objects, but tracking the storm centroid circumvents this issue. However, accounting for some of the variability of a storm by tracking the entire storm cell and any encompassing larger or neighboring storms using additional hidden variables in the graphical model could still potentially improve predictions.

VI. FUTURE WORK

There are several directions for future research. With respect to the problem of tracking meteorological phenomena, one interesting direction would be to look at multiple target-tracking: e.g., tracking the splitting and merging of storms. This would eliminate the reliance on SCIT for associating new datapoints with existing storm tracks or for identifying new storm tracks. Another interesting direction would be to incorporate more meteorological information into the SKF model, possibly by adding more hidden nodes. For instance, knowledge of the average velocity of a storm or the calibration error of the radars could be incorporated into the parameters. With respect to the models used, in this paper we focused on two related DBNs, the KF and the SKF. We are interested in exploring related models, such as an extended version of the SKF with multiple higher level layers instead of only a single switch layer. Such models would potentially permit multiscale predictions similar to those of [13], [14]. Since exact inference in models like the SKF is often computationally expensive, exploring approximate inference methods for such models would also be useful. Finally, it would be interesting to use the predictions made by models like the KF and SKF for decision-making.

ACKNOWLEDGMENT

This work was supported in part by the National Science Foundation under grants EEC-0313747001 and ECS-0218125. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. Victoria Manfredi was supported by a National Science Foundation graduate research fellowship.

REFERENCES

- [1] Latitude and longitude. <http://core.ecu.edu/geology/woods/Lat-Long.htm>.
- [2] Y. Bar-Shalom. *Estimation and tracking : principles, techniques, and software*. Artech House, Boston, Massachusetts, 1993.
- [3] R. Brooks, P. Ramanathan, and A. M. Sayeed. Distributed target classification and tracking in sensor networks. *Proceedings of the IEEE*, 91:8:1163–1171, 2003.
- [4] M. Dixon and G. Weiner. TITAN: Thunderstorm identification, tracking analysis and nowcasting a radar based methodology. *J. Atmos. Ocean. Tech.*, 10:785–797, 1993.
- [5] G. Evensen. The ensemble Kalman filter: Theoretical formulation and practical implementation. *Ocean Dynamics*, 53:343–367, 2003.
- [6] Z. Ghahramani and G. Hinton. Variational learning for switching state-space models. *Neural Computation*, 12:831–864, 2000.
- [7] Kurt Hondl. Capabilities and components of the warning decision and support system - integrated information (wdss-ii). In *Proc. American Meteorological Society Annual Meeting*, January 2003.
- [8] Bryan Horling, Roger Mailler, and Victor Lesser. A Case Study of Organizational Effects in a Distributed Sensor Network. In *Proceedings of the International Conference on Intelligent Agent Technology (IAT 2004)*, September 2004.
- [9] P. Houtekamer and H. Mitchell. Data assimilation using an ensemble Kalman filter technique. *Monthly Weather Review*, 126:796–811, 1998.
- [10] J. Johnson, P. MacKeen, A. Witt, E. DeWayne Mitchell, G. Stumpf, M. Eilts, and K. Thomas. The storm cell identification and tracking algorithm: An enhanced WSR-88D algorithm. *Weather and Forecasting*, 13:263–276, 1998.
- [11] T. Kailath, A. H. Sayed, and B. Hassibi. *Linear Estimation*. Prentice Hall, Upper Saddle River, New Jersey, 2000.
- [12] C-J. Kim. Dynamic linear models with Markov-switching. *Journal of Econometrics*, 60:1–22, 1994.
- [13] V. Lakshmanan. *A hierarchical, multiscale texture segmentation algorithm for real-world scenes*. PhD thesis, University of Oklahoma, 2001.
- [14] V. Lakshmanan, R. Rabin, and V. DeBrunner. Multiscale storm identification and forecast. *Journal of Atmospheric Research*, pages 367–380, 2003.
- [15] V. Lesser, C. Ortiz, and M. Tambe editors. *Distributed Sensor Networks: A multiagent perspective*. Kluwer Academic Publishers, Netherlands, 2003.
- [16] J. Liu, J. Liu, J. Reich, P. Cheung, and F. Zhao. Distributed group management for track initiation and maintenance in target localization applications. *Proceedings of the Second International Workshop on Information Processing in Sensor Networks (IPSN '03)*, pages 113–128, 2003.
- [17] V. Manfredi, S. Mahadevan, and J. Kurose. Kalman filters for prediction and tracking in an adaptive sensor network. *Technical report, U. of Massachusetts Amherst*, 2005.
- [18] H. Mitchell and P. Houtekamer. An adaptive ensemble Kalman filter. *Monthly Weather Review*, 126:416–433, 1998.
- [19] K. Murphy. Switching Kalman filters. *Technical report, U. C. Berkeley*, 1998.
- [20] K. Murphy. *Dynamic Bayesian Networks: Representation, Inference and Learning*. PhD thesis, University of California, Berkeley, Computer Science Division, 2002.
- [21] N. Nguyen, H. Bui, S. Venkatesh, and G. West. Multiple camera coordination in a surveillance system. *ACTA Automatica Sinica*, 29:408–422, 2003.
- [22] E. Ott, B. Hunt, I. Szunyogh, A. Zimin, E. Kostelich, M. Corazza, E. Kalnay, D. Patil, and J. Yorke. A local ensemble Kalman filter for atmospheric data assimilation. *Tellus*, 2004.
- [23] P. V. Pahalawatta, D. Depalov, T. N. Pappas, and A. K. Katsagelos. Detection, classification, and collaborative tracking

- of multiple targets using video sensors. *Proceedings of the Second International Workshop on Information Processing in Sensor Networks (IPSN '03)*, pages 529–544, 2003.
- [24] V. Pavlovic, J. Rehg, T-J. Cham, and K. Murphy. A dynamic bayesian network approach to figure tracking using learned dynamic models. In *Intl. Conf on Computer Vision*, 1999.
- [25] C. E. Pierce, P. J. Hardaker, C. G. Collier, and C. M. Haggett. GANDOLF: a system for generating automated nowcasts of convective precipitation. *Meteorol. Appl.*, 7:341–360, 2000.
- [26] B. Ristic, S. Arulampalam, and N. Gordon. *Beyond the Kalman Filter: Particle Filters for Tracking Applications*. Artech House, Boston, Massachusetts, 2004.
- [27] S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach, 2nd Edition*. Prentice Hall, Upper Saddle River, New Jersey, 2003.
- [28] G. Strang. *Introduction to Linear Algebra, 3rd Edition*. Wellesley-Cambridge Press, Wellesley, Massachusetts, 2003.
- [29] M. Wolfson, B. Forman, R. Hallowell, and M. Moore. The growth and decay storm tracker. In *American Meteorological Society 79th Annual Conference*, 1999.
- [30] B. Xu, C.K. Wikle, and N.I. Fox. A kernel-based spatio-temporal dynamical model for nowcasting radar precipitation. *Submitted to the J. Amer. Stat. Assoc.*, 2003.
- [31] M. Zink, D. Westbrook, S. Abdallah, B. Horling, V. Lakamraju, E. Lyons, V. Manfredi, J. Kurose, and K. Hondl. Meteorological command and control: An end-to-end architecture for a hazardous weather detection sensor network. *Workshop on End-to-End, Sense-and-Respond Systems, Applications, and Services (EESR)*, 2005.