IDENTIFYING AND TRACKING TURBULENCE STRUCTURES*

Curtis Storlie,[†] *Chris Davis,*[‡] *Timothy Hoar,*[§] *Thomas Lee,*[¶]
Douglas Nychka,[‖] Jeffrey B. Weiss[**] & Brandon Whitcher[††]

ABSTRACT

We present a statistical approach to object tracking which allows for paths to merge together or split apart. Paths are also allowed to be born, die, and go undetected for several frames. The splitting and merging of paths is a novel addition for a statistically-based tracking algorithm. This addition is essential for storm tracking, which is the motivation for this work. The utility of this tracker extends well beyond the tracking of storms however. It can be valuable in other tracking applications that have splitting or merging, such as vortices, radar/sonar signals, or groups of people. The method assumes that the location of an object behaves like a Gaussian Process when it is observable. Objects are required to be born, die, split, or merge according to a Markov State Model. Path correspondence is achieved by an algorithm that finds the paths that maximize the likelihood of the assumed model.

1. INTRODUCTION

The problem of object tracking has many different names, including the correspondence problem, motion correspondence, feature point tracking, and data association among others. It has applications in radar and signal processing, air traffic control, robot vision, GPS-based navigation, and biomedical engineering, to name a few. The approach developed in this paper is motivated by the need to track turbulence structures.

1.1. Description of the Tracking Problem

The basic premise of the object tracking problem is that we wish to follow objects of interest through a sequence of images. To illustrate the idea further, assume there are M objects at each time step (we are ignoring birth and death of the objects for now) as in Figure 1. The objects in the figure are labeled accoring to time so that time is distinguishable when all times are included in the plots at the bottom of the figure. The objects are free to move and change attributes from one time step (image) to the next. The purpose of the tracking algorithm is to determine which temporal set of locations corresponds to one particular object.

Tracking algorithms are dependent on an identification algorithm. The identification algorithm will go through each image and record the location of each object that it finds. In some cases, such as in ours, it may record other information about the object as well, such as size or shape.

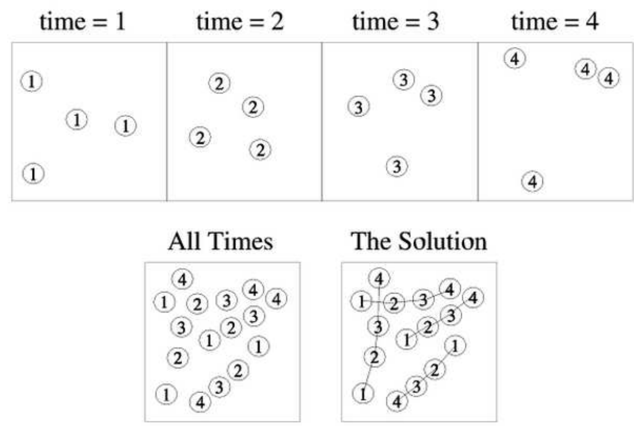

Fig. 1. A tracking problem with no birth or death.

The goal of the tracking algorithm then is to take this location and time data and assign labels to every object found by the identification algorithm in each image. In the end, we wish to recover the path of each object; that is, the history of the (x,y) coordinates of each object.

Of course, it isn't quite this easy in real life. Most applications will not have the same number of objects in each image. Imperfect detection and occlusion will lead to missing objects in some images. There can also be spurious observations or false alarms (often called clutter). In addition, some objects may appear for the first time or disappear for good in the middle of the image sequence. These events, which we will call birth and death, can happen when an object moves into or out of the field of vision. In the case of storm tracking, the objects can actually form or disappear right in the middle of the field of vision. In the case of turbulence, like-spinning structures often merge together.

1.2. Two Scientific Problems

This subsection describes two scientific problems for which the tracking of objects is an important step to their solutions. The first problem concerns the study of the evolution of storm/rainfall systems captured by satellite radar imaging techniques; see Figure 2. The very short term behavior (less than 1 hour) of such systems are reasonably well-known, but the short term (1 to 6 hours) and long term (1 to 2 days) behaviors are still largely unknown. It is definitely desirable if such longer term behaviors are better understood, and a useful tool that would help in this direction is a procedure that automatically monitors the movements and the interactions (e.g., merging and splitting) of all the structures in the overall full system. The second motivating problem focuses on a particular type of simulated image sequences of freely decaying vortices; see Figure 3. Recently such image sequences are a subject of much research [1, 8, 10], as it is (i) a paradigm for

*THIS WORK IS SUPPORTED BY NSF GRANT DMS-9815344 AND THE NATIONAL CENTER FOR ATMOSPHERIC RESEARCH

†Department of Statistics, Colorado State University
‡MMM, National Center for Atmospheric Research
§GSP, National Center for Atmospheric Research
¶Department of Statistics, Colorado State University
‖GSP, National Center for Atmospheric Research
**PAOS, University of Colorado at Boulder
††GlaxoSmithKline

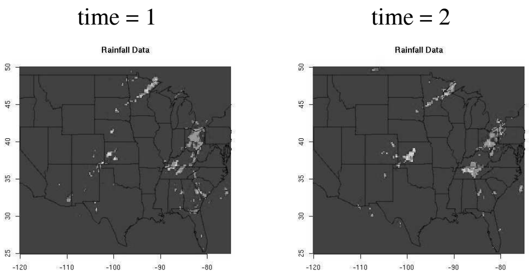

Fig. 2. Radar images of storm/rainfall.

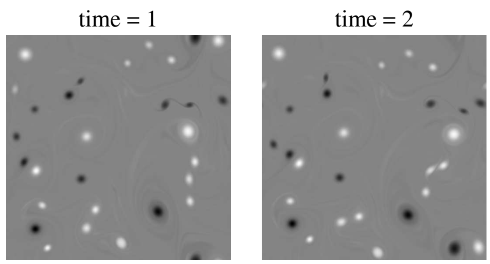

Fig. 3. Simulated images of freely decaying vortices.

anisotropic geophysical and astrophysical turbulence, and (ii) it is also the most computationally accessible example of fluid turbulence. Automatic tracking of such vortices is a first step to the understanding of the interactions amongst these structures.

A common characteristic to these two tracking problems is that the objects of interests exhibit merging, and for the former problem, the objects also exhibit splitting. Thus many existing tracking methods that do not allow merging or splitting (e.g., [4, 5, 7, 9]) can not be applied. Other methods that do attempt to incorperate merging and splitting such as [2, 3, 6] are ad-hoc as they do not take advantage of the power that comes from using a statistical model. In Section 2 we present a novel statistical tracking model that is designed to handle merging and splitting. One attractive property of this model is that its likelihood function can be quickly and accurately approximated; this is described in Section 3 below. In Section 4 we present some simulated results, while future work is discussed in Section 5.

2. A STATISTICAL MODEL

Define a path to be $\{[X(t), Y(t)] : t > 0\}$, which is the x, y coordinates of an object at all times $t > 0$. We observe the path at discrete times $\boldsymbol{T} = (T_0, T_1, \ldots, T_n)$. We wish to model the path of an object (for example a storm) by a 2 dimensional Integrated Brownian Motion. The problem is that we may not observe the process if we look at time t for several reasons: (1) It may not be found by our detection procedure, (2) it no longer exists, (3) it merged with another path, (4) it split off into 2 new paths, or (5) it may not yet exist. In addition, there may be spurious paths (false alarms) found by the detection algorithm. In the following, we present a model for the objects which takes all of these things into account. The proposed solution of the tracking problem is then the set of paths that maximize the likelihood of this model.

2.1. State Model

The *State Model* is a continuous time Markov chain that determines when paths comes into existence and terminate. Paths can either be born, die, split into two paths, or merge with another path. The rate at which these events happen are λ_1, $N(t)\lambda_2$, $N(t)\lambda_3$, and $\binom{N(t)}{2}\lambda_4$ respectively, where $N(t)$ is the number of paths in existence at time t.

This process also allows for false alarm paths to be born and die. A false alarm does not interact with the true paths or other false alarms at all (no merging or splitting). False alarms will be born with rate ρ_1 and die with rate $N_f(t)\rho_2$ where $N_f(t)$ is the number of false alarms in existence at time t. The initial number of paths and false alarms, $N(0)$ and $N_f(0)$ are assumed to be Poisson distributed with rates λ_0 and ρ_0 respectively.

The following notation will be used to describe the states of each path. Let each path that will exist before the last image at time T_n be indexed by $i = 1, \ldots, M$. Let p_i denote the two parents of the i^{th} path. If path i results from a birth, then $p_i = (0, 0)$ represents that it has no parents. For a path resulting from a split of path k, $p_i = (k, 0)$. If i is a path resulting from a merger of paths k and l, then $p_i = (k, l)$. Finally, let $p_i = (-1, -1)$ to represent a false alarm.

Now let ξ_i denote the initiation time of path i and ζ_i the time of termination. So the collection of $\boldsymbol{p} = (p_1, \ldots, p_M)$, $\boldsymbol{\xi} = (\xi_1, \ldots, \xi_M)$ and $\boldsymbol{\zeta} = (\zeta_1, \ldots, \zeta_M)$ contains all of the information in the State Model.

2.2. Missing State Model

We now deal with the problem where a path exists but is not observable for some reason. This behavior will be modelled with another continuous time Markov Chain with state variable $W(t)$ that simply takes on the values 0 if missing and 1 if observable. When missing, the path becomes observable with rate ν_0 and when observable, the path becomes missing with rate ν_1.

Let $W_i(t)$ denote the missing state variable of the i^{th} path at time t. Let $W_i(t) = 0$ for the i^{th} path until it is initiated. As soon as path i is initiated, the value of $W_i(\xi_i)$ is determined by the steady state probabilities of the chain. As soon as path i is terminated by the State Model, $W_i(t)$ becomes 0 again for $t \geq \zeta_i$.

False alarm paths are assumed to always be observable when they exist. So we can assume $\nu_0 = 1$ and $\nu_1 = 0$ for example. This will ensure that a false alarm is never missing when it exists.

2.3. Object Size

The minor and major axes, R_1 and R_2 respectively, of the best fitting ellipse to each object will also be in the model (Figure 4). The best fitting ellipse is determined by standard imaging techniques. We can also define the size to be $S = R_1 R_2$.

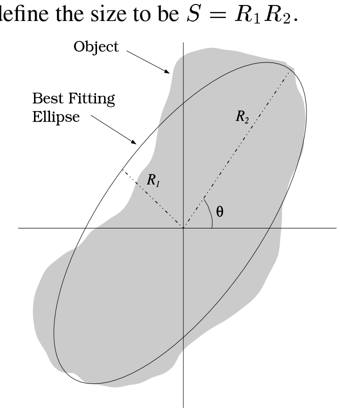

Fig. 4. R_1, R_2, and θ

Let $R_{1,i}(t)$ and $R_{2,i}(t)$ be the radii of the minor and major axes respectively of path i at time t. $R_{1,i}(t)$ and $R_{2,i}(t)$ will be treated as random variables from a lognormal distribution with parameters $(\mu_{R_{1,i}}, \sigma^2_{R_{1,i}})$ and $(\mu_{R_{2,i}}, \sigma^2_{R_{2,i}})$. These observations are assumed to be independent over time.

Since $R_{1,i}(t)$ and $R_{2,i}(t)$ are lognormal, $S_i(t)$ is also lognormal with parameters $(\mu_{R_1,i} + \mu_{R_2,i}, \sigma^2_{R_1,i} + \sigma^2_{R_2,i})$. In the event of a split, the mean sizes of the two new objects are required to sum to the mean size of the parent object. If ES_i is the mean size of the i^{th} path which splits into paths j and k, then we have $ES_i = ES_j + ES_k$. In the event of a merger, the same conservation of mean size rule will apply.

In reality, we only observe the order statistics $R_{(1)}(t) = R_1(t) \wedge R_2(t)$ and $R_{(2)}(t) = R_1(t) \vee R_2(t)$. Thus it is the likelihood of $R_{(1)}$ and $R_{(2)}$ that needs to be calculated in Section 3.

2.4. Object Orientation

The orientation of an object will be measured by the angle, $\theta(t)$ of the major axis, $R_2(t)$ at time t as shown in Figure 4. We will use the usual notation $\theta_i(t)$ to represent the angle for the i^{th} path. We will assume that $\theta_i(t)$ comes from a VonMises distribution on $[0, \pi]$ with parameters α_i and β_i and *iid* over time.

Since $R_2(t)$ is unobservable, neither is $\theta(t)$. We do observe the angle of $R_{(2)}(t)$ which we will call $\theta'(t)$. The likelihood for $\boldsymbol{\theta}'$ will be given in Section 3.

2.5. Object Location

Let the x-coordinate of the i^{th} path at time t be denoted by $X_i(t)$. The distribution of $X_i(t)$ will be defined below. Therefore $Y_i(t)$ will be the same with the obvious changes in notation and independent of $X_i(t)$. Suppose that the i^{th} path resulted from a birth. Then,

$$X_i(t) = X_i(\xi_i) + X_i'(\xi_i)(t - \xi_i) + \sigma_i Z_i(t - \xi_i), \quad (1)$$

where $X_i'(t)$ is the velocity of the path at time t and $Z_i(t)$ is Integrated Brownian Motion given by, $Z_i(t) = \int_0^t B_i(s)ds$, where $B_i(t)$ is the Brownian Motion driving the i^{th} path. Also recall that ξ_i is the time of initiation of path i. It is assumed that the initial position and velocity are Gaussian. Specifically, $X_i(\xi_i) \sim N(\mu_{X_0}, \sigma^2_{X_0})$ and $X_i'(\xi_i) \sim N(\mu_{X_0'}, \sigma^2_{X_0'})$.

For a path resulting from a split,

$$X_i(t) = X_{p_{i,1}}(\xi_i) + \phi_i + \left[X'_{p_{i,1}}(\xi_i) + \phi_i'\right](t - \xi_i) + \sigma_i Z_i(t - \xi_i), \quad (2)$$

where $\phi_i \sim N(0, \sigma^2_{X_s})$ and $\phi_i' \sim N(0, \sigma^2_{X_s'})$. Notice that the initial position and velocity of the new path are the same as that of the parent path plus some error. It is assumed that $\sigma^2_{X_s}$ is small so that the new paths appear close to where the parent split. There will also be something like a conservation of momentum assumption imposed on the two new paths splitting off of the parent. This is achieved by the following.

Let c_i be a vector containing the three paths involved in the i^{th} splitting event, $i = 1, \ldots, N_s$, where N_s is the number of splits before time T_n. The index of the parent path is $c_{i,1}$ where $c_{i,2}$ and $c_{i,3}$ are the children. The change in momentum after the split is

$$\begin{aligned} C_i &= ES_{c_{i,2}} X'_{c_{i,2}}(\xi_{c_{i,2}}) + ES_{c_{i,3}} X'_{c_{i,3}}(\xi_{c_{i,2}}) \\ &\quad - ES_{c_{i,1}} X'_{c_{i,1}}(\xi_{c_{i,2}}). \end{aligned} \quad (3)$$

We will then condition the model for $\boldsymbol{X}$ on the events $C_i = 0$ for $i = 1, \ldots, N_s$. This will ensure that there is no change in momentum.

For a path resulting from a merging,

$$\begin{aligned} X_i(t) &= \frac{ES_{p_{i,1}}}{ES_i} X_{p_{i,1}}(\xi_i) + \frac{ES_{p_{i,2}}}{ES_i} X_{p_{i,2}}(\xi_i) + \\ &\left[\frac{ES_{p_{i,1}}}{ES_i} X'_{p_{i,1}}(\xi_i) + \frac{ES_{p_{i,2}}}{ES_i} X'_{p_{i,2}}(\xi_i)\right](t - \xi_i) + \sigma_i Z_i(t - \xi_i). \end{aligned} \quad (4)$$

Ensuring that the two parent paths get close together before merger will be accomplished in a similar manner to the conservation of momentum for splitting.

Let d_i be a vector containing the three paths involved in the i^{th} merger, $i = 1, \ldots, N_m$, where N_m is the number of mergers before time T_n. The indices of the parent paths are $d_{i,1}$ and $d_{i,2}$, where $d_{i,3}$ is the index of the child. The difference in location between the two parents at the time of merger plus a small error is given by

$$D_i = X_{d_{i,1}}(\xi_{d_{i,3}}) - X_{d_{i,2}}(\xi_{d_{i,3}}) + \psi_i, \quad (6)$$

where $\psi_i \sim N(0, \sigma^2_{X_m})$. We will then condition the model for X on the event $D_i = 0$ for $i = 1, \ldots, N_m$. If σ_{X_m} is small, it will force the two parent paths close together right before the merger.

Lastly in the case that the i^{th} path is a false alarm, we assume a Brownian Motion model,

$$X_i(t) = X_i(\xi_i) + \sigma_i B_i(t - \xi_i). \quad (7)$$

3. MODEL LIKELIHOOD

In this Section we will present the likelihood of the model described in Section 2. We wish to write out the density, for the collection of random variables, $\Phi = (\boldsymbol{p}, \boldsymbol{\xi}, \boldsymbol{\zeta}, \boldsymbol{W}, \boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{R}_{(1)}, \boldsymbol{R}_{(2)}, \boldsymbol{\theta}')$, where the bold variables denote the collection of values for all paths at all times. Using the property $f_{X,Y} = f_X f_{Y|X}$, we can write this density

$$\begin{aligned} f_\Phi(\phi) &= f_{\boldsymbol{p},\boldsymbol{\xi},\boldsymbol{\zeta},\boldsymbol{W},\boldsymbol{X},\boldsymbol{Y},\boldsymbol{R}_{(1)},\boldsymbol{R}_{(2)},\boldsymbol{\theta}'} \\ &= f_{\boldsymbol{p},\boldsymbol{\xi},\boldsymbol{\zeta}} \cdot f_{\boldsymbol{W}|\boldsymbol{p},\boldsymbol{\xi},\boldsymbol{\zeta}} \cdot f_{\boldsymbol{X}|\boldsymbol{p},\boldsymbol{\xi},\boldsymbol{\zeta},\boldsymbol{W}} \cdot f_{\boldsymbol{Y}|\boldsymbol{p},\boldsymbol{\xi},\boldsymbol{\zeta},\boldsymbol{W}} \cdot \\ &\quad f_{\boldsymbol{R}_{(1)},\boldsymbol{R}_{(2)}|\boldsymbol{p},\boldsymbol{\xi},\boldsymbol{\zeta},\boldsymbol{W}} \cdot f_{\boldsymbol{\theta}'|\boldsymbol{p},\boldsymbol{\xi},\boldsymbol{\zeta},\boldsymbol{W},\boldsymbol{R}_{(1)},\boldsymbol{R}_{(2)}}. \end{aligned} \quad (8)$$

since $\boldsymbol{X}$, $\boldsymbol{Y}$, and $(\boldsymbol{R}_{(1)}, \boldsymbol{R}_{(2)})$ are independent given $(\boldsymbol{p}, \boldsymbol{\xi}, \boldsymbol{\zeta}, \boldsymbol{W})$. Also $\boldsymbol{\theta}'$ is independent of $\boldsymbol{X}$ and $\boldsymbol{Y}$ given $(\boldsymbol{p}, \boldsymbol{\xi}, \boldsymbol{\zeta}, \boldsymbol{W}, \boldsymbol{R}_{(1)}, \boldsymbol{R}_{(2)})$. We will call the conditional densities in (8), in order from left to right, the state density, missing state density, X density, Y density, radius density, and angle density respectively.

3.1. State Density

The state variables $\boldsymbol{p}$, $\boldsymbol{\xi}$, and $\boldsymbol{\zeta}$, represent the state model in a convenient form, but the density for these variables is very difficult to write out. We will therefore represent the state model with some equivalent variables for the purpose of writing out a likelihood.

Let τ_i be the time of the i^{th} event in the State Model, $i = 1, \ldots, N_e$, where N_e is the number of events before time T_n. An event is considered to be a birth, death, split, merger, false birth, or false death. Let U_i represent the i^{th} event in the State Model, so that $U_i = 1$ for a birth, 2 for a death, 3 for a split, 4 for a merger, 5 for a false birth, and 6 for a false death. Let V_i contain the indices of the path(s) involved in the i^{th} event.

Recall that $N(0)$ and $N_f(0)$ are the initial number of true paths and false alarms respectively. So $\boldsymbol{\tau} = (\tau_1, \ldots, \tau_{N_e})$, $\boldsymbol{U} = (U_1, \ldots, U_{N_e})$, $\boldsymbol{V} = (V_1, \ldots, V_{N_e})$, $N(0)$, and $N_f(0)$ are an equivalent set of variables to describe the state model, in the sense that there is a 1-1 transformation from these variables to $(\boldsymbol{p}, \boldsymbol{\xi}, \boldsymbol{\zeta})$.

We do not actually get to observe these states from the data. Determining the correct states (which path merged with which etc.), is part of finding the paths that maximize the likelihood.

Now, returning to the density of the State Model, using properties of conditional probability, we have

$$f_{N(0),N_f(0),\boldsymbol{U},\boldsymbol{V},\boldsymbol{\tau}} = f_{N(0)} f_{N_f(0)} f_{\boldsymbol{U},\boldsymbol{V},\boldsymbol{\tau}|N(0),N_f(0)}, \tag{9}$$

where $N(0)$ and $N_f(0)$ are Poisson with parameters λ_0 and ρ_0, so

$$f_{N(0)}(k) = \frac{\lambda_0^k e^{-\lambda_0}}{k!}, \qquad f_{N_f(0)}(k) = \frac{\rho_0^k e^{-\rho_0}}{k!}. \tag{10}$$

Now suppose we use the convention that $I_A(x)$ is the indicator function for $x \in A$. Also write $\lambda_1^*(t) = \lambda_1$, $\lambda_2^*(t) = N(t)\lambda_2$, $\lambda_3^*(t) = N(t)\lambda_3$, $\lambda_4^*(t) = \binom{N(t)}{2}\lambda_4$, $\lambda_5^*(t) = \rho_1$, $\lambda_6^*(t) = N_f(t)\rho_2$, and $\lambda^*(t) = \sum_{i=1}^6 \lambda_i^*(t)$. The conditional density on the right side of (9) can then be written

$$f_{\boldsymbol{U},\boldsymbol{V},\boldsymbol{\tau}|N(0),N_f(0)}(u,v,t) = \prod_{i=1}^{N_e} \Bigg\{ \left(\lambda^* e^{-\lambda^*(t_i-\tau_{i-1})}\right) \left(\frac{\lambda^*_{u_i}}{\lambda^*}\right)$$
$$\left(I_{\{1,5\}}(u_i) + \frac{I_{\{2,3\}}(u_i)}{N(\tau_{i-1})} + \frac{I_{\{4\}}(u_i)}{\binom{N(\tau_{i-1})}{2}} + \frac{I_{\{6\}}(u_i)}{N_f(\tau_{i-1})} \right) \Bigg\}. \tag{11}$$

3.2. Missing State Density

Let $W_{i,j} = W_i(T_j)$. Since for $i \neq j$, W_i is independent of W_j given the state variables and W_i is Markov, we have

$$f_{\boldsymbol{W}|\boldsymbol{p},\boldsymbol{\xi},\boldsymbol{\zeta}}(w) = \prod_{i=1}^{M}\prod_{j=1}^{n} f_{W_{i,j}|\boldsymbol{p},\boldsymbol{\xi},\boldsymbol{\zeta},W_{i,j-1}}(w_j). \tag{12}$$

Now let $P_{j,k}(t)$ be the probability of W_i going from state j to state k in a time t , assuming that the path exists during this time. These transition probabilities are

$$P_{j,k}(t) = \frac{\nu_{1-k}}{\nu_0+\nu_1} + \frac{\nu_j}{\nu_0+\nu_1} e^{-(\nu_0+\nu_1)t}. \tag{13}$$

The stationary distribution of the chain is

$$\pi_k = \lim_{t\to\infty} P_{j,k}(t) = \frac{\nu_{1-k}}{\nu_0+\nu_1}. \tag{14}$$

Also let $\Delta T_j = T_j - T_{j-1}$.

We can then write the conditional density of $W_{i,j}$ on the right side of (12) by using indicators to break it apart into the times when the path exists and does not. In some instances we use the indicator I_A without the argument (x). In that case use the convention that I_A equals 1 if the event A occurred and 0 otherwise:

$$f_{W_{i,j}|\boldsymbol{p},\boldsymbol{\xi},\boldsymbol{\zeta},W_{i,j-1}}(w_j) = I_{\{T_j<\xi_i\}} I_{\{T_j>\zeta_i\}} I_{\{0\}}(w_j) + \tag{15}$$
$$I_{\{T_{j-1}<\xi_i\}} I_{\{\xi_i\le T_j\le\zeta_i\}} \pi_{w_j} + I_{\{T_{j-1}>\xi_i\}} I_{\{T_j<\zeta_i\}} P_{W_{j-1},w_j}(\Delta T_j).$$

3.3. Radius Density

Recall that $R_1(t)$ and $R_2(t)$ are distributed as independent lognormals and we observe the min and max of these which are $R_{(1)}(t)$ and $R_{(2)}(t)$ respectively. The density for $(R_{(1)}(t), R_{(2)}(t))$ is similar to that for order statistics

$$f_{R_{(1)},R_{(2)}}(r,s) = [f_{R_1}(r)f_{R_2}(s) + f_{R_1}(s)f_{R_2}(r)]\, I_{\{r\le s\}}. \tag{16}$$

where f_{R_1} and f_{R_2} are lognormal densities as described in Section 2.3

The density of the radii in (8) is conditional on p, ξ, ζ, and W, but this density really only depends on W so we will drop the other subscripts in the following. Since the radii of path i at time T_j, $(R_{(1),i,j}, R_{(2),i,j})$, are independent of the radii at other times and paths i and j are independent, the density in (8) can be written

$$f_{\boldsymbol{R}_{(1),i},\boldsymbol{R}_{(2),i}|\boldsymbol{W}}(r,s) = \prod_{i=1}^{M} \prod_{\{j:W_{i,j}=1\}} f_{R_{(1),i},R_{(2),i}}(r_{i,j}, s_{i,j}). \tag{17}$$

3.4. Angle Density

The distribution of $\theta'(t)$ given $(R_{(1)}(t), R_{(2)}(t))$ is a mixture distribution that takes the value of $\theta(t)$ with probability γ and $\lfloor \theta(t) + \pi/2 \rfloor_\pi$ with probability $1-\gamma$, where $\lfloor x \rfloor_y$ is x mod y and

$$\begin{aligned} \gamma &= P(R_1(t) < R_2(t) \mid R_{(1)}(t) = r, R_{(2)}(t) = s) \\ &= \frac{f_{R_1}(r)f_{R_2}(s)}{f_{R_1}(r)f_{R_2}(s) + f_{R_1}(s)f_{R_2}(r)}. \end{aligned} \tag{18}$$

Hence the conditional density of $\theta'(t)$ is

$$f_{\theta'(t)|R_{(1)}(t),R_{(2)}(t)}(z) = \gamma f_{\theta(t)}(z) + (1-\gamma) f_{\theta(t)}(\lfloor z+\pi/2 \rfloor_\pi), \tag{19}$$

where $f_{\theta(t)}$ is the VonMises density on $[0,\pi)$. As with the radii, θ_i is independent over time and of other paths. Let $\theta'_{i,j} = \theta'_i(T_j)$ and we have

$$f_{\theta'|W,R_{(1)},R_{(2)}}(z) = \prod_{i=1}^{M} \prod_{\{j:W_{i,j}=1\}} f_{\theta'_{i,j}|R_{(1),i,j},R_{(2),i,j}}(z_{i,j}). \tag{20}$$

3.5. Density of X

Since $X_i(t)$ is normally distributed for all t, the collection of locations of all paths at all observed time points, $\boldsymbol{X}$ has a multivariate normal distribution, $\boldsymbol{X} \sim N(\boldsymbol{\mu}_X, \Sigma_X)$. Recall in section 2.5 that we need to then condition $\boldsymbol{X}$ on $C = 0$ and $D = 0$ for $C = (C_1, \ldots, C_{N_s})$ and $D = (D_1, \ldots, D_{N_m})$, which are also normally distributed, $C \sim N(\boldsymbol{\mu}_C, \Sigma_C)$, $D \sim N(\boldsymbol{\mu}_D, \Sigma_D)$. For the collection of all three vectors we have $(\boldsymbol{X}, C, D) \sim N(\boldsymbol{\mu}, \Sigma)$ where $\boldsymbol{\mu} = (\boldsymbol{\mu}_X, \boldsymbol{\mu}_C, \boldsymbol{\mu}_D)'$ and

$$\Sigma = \left(\begin{array}{c|cc} \Sigma_X & \Sigma_{X,C} & \Sigma_{X,D} \\ \hline \Sigma_{C,X} & \Sigma_C & \Sigma_{C,D} \\ \Sigma_{D,X} & \Sigma_{D,C} & \Sigma_D \end{array}\right) = \left(\begin{array}{c|c} \Sigma_X & \Lambda \\ \hline \Lambda' & \Gamma \end{array}\right).$$

Multivariate normal theory tells us that the distribution of $\boldsymbol{X}$ given $C = 0$ and $D = 0$ is $N(\boldsymbol{\mu}^*, \Sigma^*)$ where

$$\boldsymbol{\mu}^* = \boldsymbol{\mu}_X - \Lambda\Gamma^{-1} \begin{pmatrix} \boldsymbol{\mu}_C \\ \boldsymbol{\mu}_D \end{pmatrix}, \qquad \Sigma^* = \Sigma_X - \Lambda\Gamma^{-1}\Lambda'. \tag{21}$$

The density of $\boldsymbol{X}$ conditional on $C = 0$ and $D = 0$ is then just the multivariate normal density with parameters, $\boldsymbol{\mu}^*$ and Σ^*.

4. SIMULATED EXAMPLES

Here we give two simple examples to show that maximizing the likelihood of the model described above is a reasonable way to find the tracking solution. In Figure 5 we have two paths that merge together, then split apart. The individual images at times 1-6 are given on the left while some hypotheses for the complete paths are given on the right with their corresponding log-likelihood value. The top right image is the correct hypothesis and you can see that it has a higher likelihood than any of the alternatives.

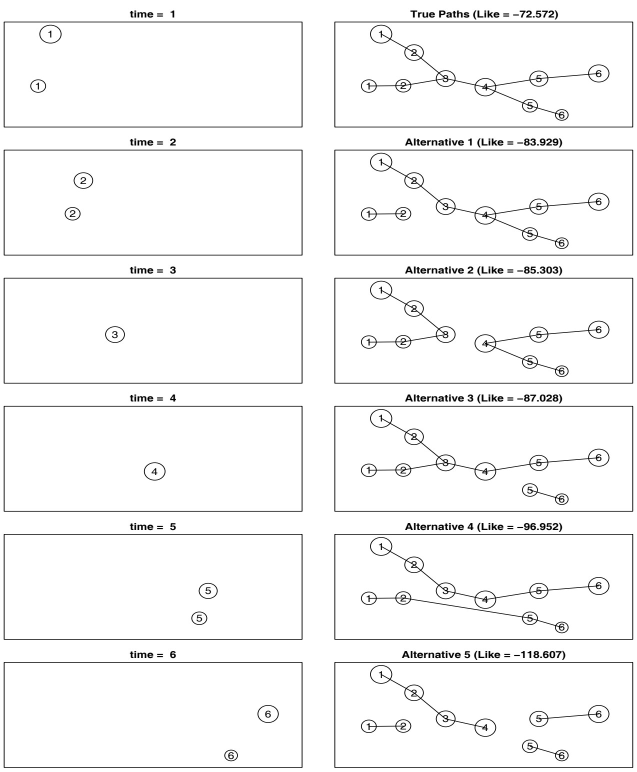

Fig. 5. A merging and splitting event with various alternative hypotheses

An example of a crossing event is given in Figure 6 to demonstrate this model's ability to distinguish between crossing and merging/splitting. Once again, the true hypothesis in the top right has the highest likelihood.

5. CONCLUSIONS

In this paper we presented a novel statistical tracking model that allows for object merging and splitting. Based on this model, a practical tracking algorithm is currently being developed. The ultimate goal is to be able to track the objects for the storm/rainfall and freely decaying vortex problems.

6. REFERENCES

[1] A. Bracco, J. C. McWilliams, G. Murante, A. Provenzale, and J. B. Weiss. Revisiting freely decaying two-dimensional turbulence at millennial resolution. *Physics of Fluids*, 12:2931–2941, 2000.

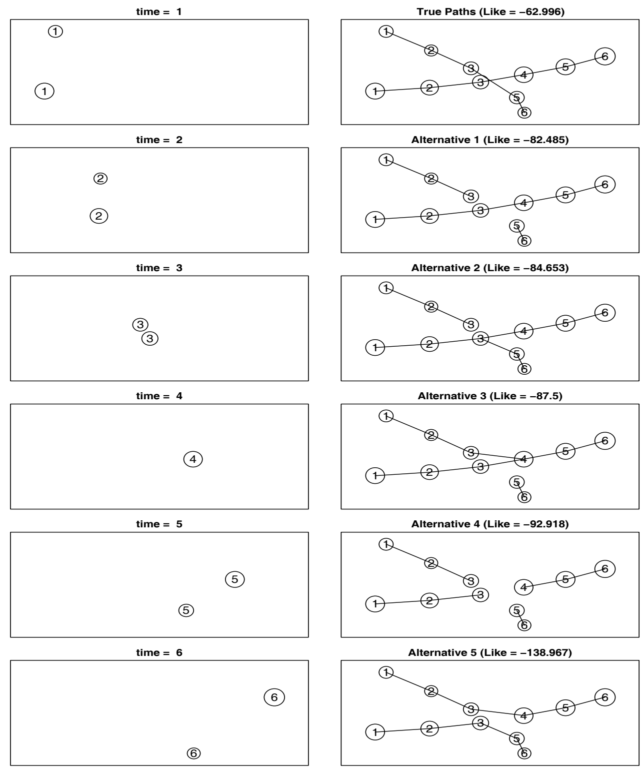

Fig. 6. Two paths crossing with various alternative hypotheses

[2] L. Carvalho and C. Jones. A satellite method to identify structural properties of mesoscale convective systems based on the maximum spatial correlation tracking technique (MASCOTTE). *Journal of Applied Meteorology*, 40:1683–1701, 2001.

[3] D. Cheng, R. Mercer, J. Barron, and P. Joe. Tracking severe weather storms in dopplar radar images. *International Journal of Imaging Systems and Technology*, 9:201–213, 1998.

[4] I. Cox. A review of statistical data association techniques for motion correspondence. *International Journal of Computer Vision*, 10:53–66, 1993.

[5] T. Fortmann, Y. Bar-Shalom, and M. Scheffe. Sonar tracking of multiple targets using joint probabilistic data association. *IEEE Journal of Oceanic Engineering*, 8:173–184, 1983.

[6] L. Fuentes and S. Velastin. Tracking people for automatic surveillance applications. *Lecture Notes in Computer Science*, 2652:238–245, 2003.

[7] K. Hodges. A general method for tracking analysis and its application to meteorological data. *Monthly Weather Review*, 122:2573–2586, 1994.

[8] C. Pasquero, A. Provenazale, and J. B. Weiss. Vortex statistics from eulerian and lagrangian time series. *Physical Review Letters*, 89:284501, 2002.

[9] D. Reid. An algorithm for tracking multiple targets. *IEEE Transactions on Automatic Control*, 24:843–854, 1979.

[10] J. B. Weiss and J. C. McWilliams. Temporal scaling behavior of decaying two-dimensional turbulence. *Physics of Fluids A: Fluid Dynamics*, 5:608–621, 1993.