Tracking with ranked signals

Tianyang Li[†] **Harsh Pareek**[†] **Pradeep Ravikumar**[†] **Dhruv Balwada**[*] **Kevin Speer**[*]

[†] Department of Computer Science, The University of Texas at Austin

{lty,harshp,pradeepr}@cs.utexas.edu

[*] Geophysical Fluid Dynamics Institute at Florida State University

{db10d,kspeer}@fsu.edu

Abstract

We present a novel graphical model approach for a problem not previously considered in the machine learning literature: that of tracking with ranked signals. The problem consists of tracking a single target given observations about the target that consist of ranked continuous signals, from unlabeled sources in a cluttered environment. We introduce appropriate factors to handle the imposed ordering assumption, and also incorporate various systematic errors that can arise in this problem, particularly clutter or noise signals as well as missing signals. We show that inference in the obtained graphical model can be simplified by adding bipartite structures with appropriate factors. We apply a hybrid approach consisting of belief propagation and particle filtering in this mixed graphical model for inference and validate the approach on simulated data. We were motivated to formalize and study this problem by a key task in Oceanography, that of tracking the motion of RAFOS ocean floats, using range measurements sent from a set of fixed beacons, but where the identities of the beacons corresponding to the measurements are not known. However, unlike the usual tracking problem in artificial intelligence, there is an implicit ranking assumption among signal arrival times. Our experiments show that the proposed graphical model approach allows us to effectively leverage the problem constraints and improve tracking accuracy over baseline tracking methods yielding results similar to the ground truth hand-labeled data.

1 INTRODUCTION

We consider the problem of tracking a single target where the observations, concerning the position of the target, consist of ranked continuous signals from unlabeled sources in a cluttered environment. We allow for various types of error that may occur in these observations: (a) a recorded signal may correspond to a spurious signal instead of a true signal, and (b) a signal may be lost and never recorded. We present a novel graphical model approach for this problem. Our graphical model is a mixed or hybrid model with both discrete and continuous components. To perform probabilistic inference in this mixed graphical model, we use particle filtering for the continuous component which represents the location of the target, and belief propagation for the discrete component which represents the data association between the signals and the signals' sources.

We were motivated to study the above problem formalism by a mathematical abstraction of a key problem in Oceanography: that of tracking RAFOS floats using ranked range measurements. RAFOS floats (Rossby *et al.* , 1986; Hancock & Speer, 2013) are low cost acoustically tracked subsurface floating devices used to study ocean currents by measuring the paths taken by fluid parcels in the ocean. They also measure temperature and pressure along the way. The typical mission times for these floats are on the order of a few months to a few of years during which they don't surface, and hence it is not possible to **locate or track** these floats via satellite (GPS) position fixes. In order to solve this location or tracking problem, a moored (fixed) array of sound sources, also known as *beacons* are used. These moored sound sources or beacons produce and transmit one sound signal per tracking cycle. The floats then record the **arrival times** of these signals transmitted from the beacons. These arrival times depend on the distance from the sound source and the velocity of sound in the ocean, and thus provide information about the location of the float (since the sound sources or beacons are fixed and their locations are known). At any given instant, if distances of the float from three beacons are known, this suffices to determine the position of the float. Thus if we could identify the beacons corresponding to the received signals, we can then track the course of the float.

The caveat is that while the arrival times are stored, the beacons from which the respective signals originated is **not** known. In addition, the storage capacity of each float is limited and it only records a small number of these signal arrival times each day. These signals have a natural ordering in that the beacon signals most likely to be stored originate from the closest beacons and are received (and stored) in distance order. We are interested in inferring the identity of the beacons at each time step corresponding to the few received ordered or "ranked" signals. This is a challenging task and currently, this information is hand-labeled by oceanography researchers, with many months of effort. Our paper provides an automated solution for this interesting problem, and moreover provides a novel machine learning problem abstraction of tracking with ranked signals, which would be of interest even from a purely machine learning standpoint.

We will nonetheless anchor our discussion of the tracking with ranked signals problem to the RAFOS float tracking problem for presentational reasons. Let us consider the setup and assumptions of the above RAFOS float tracking problem in greater detail. Each float's tracking data consists of its initial and final positions, and a fixed number of earliest signal arrival times for each day. Figure 1 shows the observed signal arrival times for a particular float over the entire tracking period. There are s fixed beacons with known positions. Each float in the ocean is equipped with a receiver. Every day at a specified time, a float starts listening for sound signals transmitted by the beacons. The float stores the arrival times of the first $r(< s)$ signals it receives and a confidence value for each signal, then shuts off its receiver. Our goal is to use these arrival times to track the position of the float over time. The model we present includes the following kinds of errors that capture key characteristics of this problem:

- The arrival times are subject to noise due to environmental factors and recording equipment
- Signals from a beacon may never reach or be dropped by the receiver, we call these errors *missing values*.
- The receiver may erroneously record ambient noise as a signal from a beacon, and store that value. We call these errors *junk values* or *clutter*. Figure 1 and our analysis in the experiments section show that junk values are very common and outnumber true signals in the data.

While our focus is on the beacon association problem, it should be noted that there are additional sources of error such as those due to the Doppler effect, variations in the speed of sound due to temperature and beacon depth or clock drift in the receiver which have been investigated in prior work(Wooding *et al.* , 2005; Hancock & Speer, 2013). These would change the conditional probability distributions in the model we introduce in Figure 2 and incorporating these in our model is left as future work.

Contributions The main contributions of this work are as follows:

- Our setting differs from prior work from a machine learning standpoint due to the ordering condition imposed on the received signals, as described earlier. Our model introduces novel ranking based factors to enforce this condition, and we present an equivalent bipartite model similar in structure to that used in the data association literature (Williams & Lau, 2010) which can then be used for inference via message passing. This contribution should be valuable for other graphical model problems where ranking based factors are involved.
- RAFOS float tracking has not previously been considered from a graphical model perspective; indeed, (Hancock & Speer, 2013; Wooding *et al.* , 2005) label data by hand. Our model in Figure 2 captures several important characteristics of this problem and is also of practical importance from an application standpoint.
- We evaluate our algorithm on simulated data and show that our algorithm performs better than a baseline model which does not take the ranking of signal arrival times into account. We also present the results of our algorithm on real world RAFOS float data and demonstrate good agreement with hand-labeled data.

1.1 RELATED WORK

We now review prior work on tracking, and specifically contrast our problem with the classical tracking problem. A tracking algorithm can be viewed as computing the probability distribution of a system's state x_t given observations y_t. A common model for this is a Hidden Markov model with hidden states x_t. Inference using message passing for this model leads to the well known Kalman filtering and smoothing algorithms or the forward-backward algorithm under different distributional assumptions. Another common class of approaches to perform inference on such models are particle filtering and other MCMC-based approaches (Oh *et al.* , 2009; Chertkov *et al.* , 2010).

In the RAFOS float tracking problem, from the perspective of the float, the problem resembles a multi-target tracking or data association problem, where the beacons are the targets. In multi-target tracking, we try to track a number of moving targets over time given some noisy information about the set of observed locations at each time step, and the challenge is to link the locations over time to obtain the trajectory for each target. Classical approaches to multi-target tracking assign a probability to each possible association of targets across time steps. The inference problem $p(x_t|y_t)$ then requires summing over all possible such associations and reduces to computing the permanent of a matrix (Oh *et al.* , 2009; Chertkov *et al.* , 2010), which is

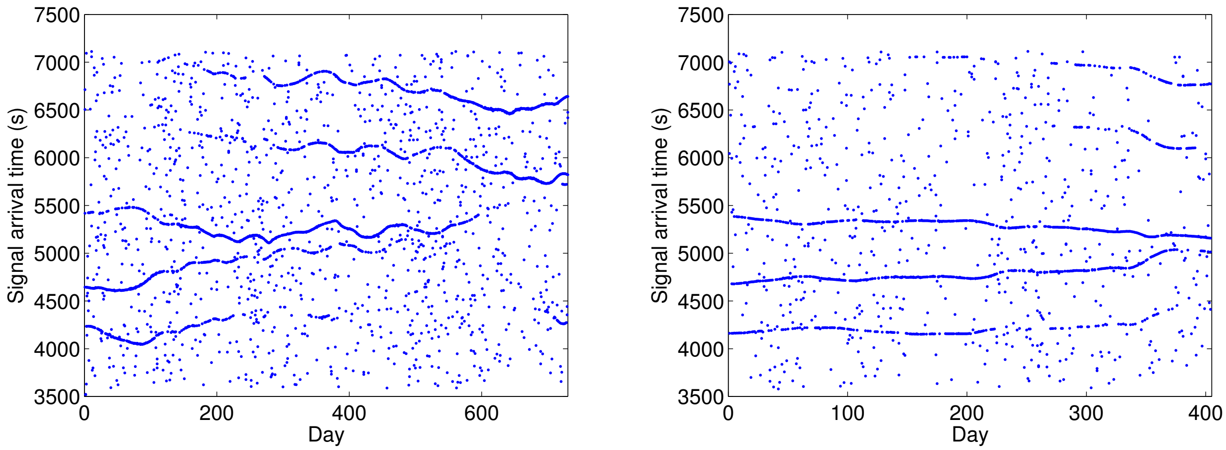

(a) float #767

(b) float #811

Figure 1: Observed signal arrival times for float #767 and #811 over the entire tracking period

known to be #P-complete (Valiant, 1979). The probabilistic data association filter (Bar-Shalom, 1987; Bar-Shalom *et al.* , 2009) collapses this state into a single Gaussian at each time step to make the problem tractable. Another interesting line of work is multiple hypothesis tracking (Blackman, 2004). At each time step the algorithm maintains a mixture of Gaussians for each target representing a distribution over its possible positions, and updates the state by summing over all possible associations while discarding components with low probability to make computation tractable. Message passing algorithms which have been shown to converge (Vontobel, 2013; Huang & Jebara, 2009) have also been proposed. (Cevher *et al.* , 2006) developed a particle filtering approach which uses only range measurements, while MCMC algorithms that are fully polynomial-time randomized approximation schemes (Jerrum *et al.* , 2004) have also been studied. Other approaches include greedy approaches widely used in robotics, which include choosing the data association with the maximum likelihood (Thrun *et al.* , 2005), and nearest-neighbor methods where observations "closest" to expected observations are kept and others are discarded. While simple, such greedy approaches work poorly under relatively high noise levels (Bar-Shalom *et al.* , 2009). Yet another interesting line of work represents the state x_t of the system using distributions over permutations and uses group-theoretic methods to approximate these distributions (Kondor *et al.* , 2007; Huang *et al.* , 2009). Such approaches have been investigated in the contexts of radar tracking, computer vision, and robotics.

We cannot directly apply these methods to the RAFOS float tracking problem because the number of observed targets (beacon arrival times) is small compared to the number of beacons and thus most beacons are unobserved at each time step. This is because if we can identify the beacons corresponding to each arrival time, only a small number (three in 2D space) of beacons are required to track the float's position, and thus the float need only store the first few arrival times while the number of beacons may be much larger. Further, there may be regime changes, where a beacon goes out of range and a previously unobserved beacon appears in range which these methods cannot directly handle. However, since the closest beacons are the ones most likely to be recorded and their values are received in order, the provided information may be sufficient for tracking. We show in this paper that explicitly modeling the ordering assumption among beacon times and the fact that there is an underlying latent variable – the float's position – which connects the targets, allows us to track the float.

We also note that multiple hypothesis tracking approaches which maintain a mixture over the potential associations are not required for our model for tracking with ranked signals. In a sense, in the RAFOS float tracking problem the only true latent variable is the float's position and the others such as the beacon arrival times are derived from this. In particular, the additional ordering assumption provides a strong helpful constraint, since the only way the closest beacon's signal is not received first is if two signals were close together and their ordering was changed due to ambient noise, if that signal is lost or if a junk signal was recorded as the first. The first possibility does not affect tracking since the two signals were already close, while the probabilities of the latter two systematic errors can be modeled in a principled manner with a graphical model approach. Thus, in the RAFOS tracking problem knowing the initial position allows us to continue tracking the float without needing a mixture model and in this sense, our model for tracking with ranked signals is (computationally and statistically) easier than the standard multiple-target tracking problem.

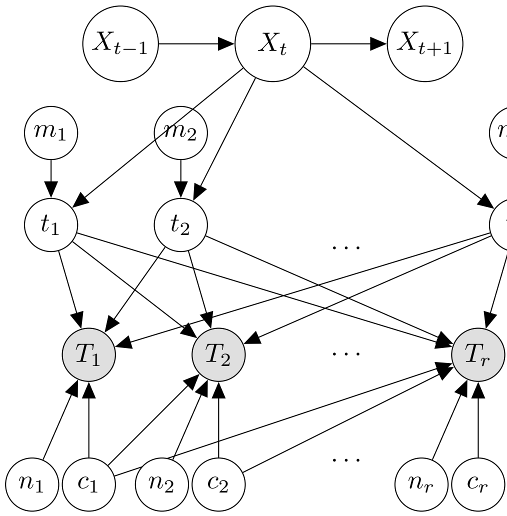

Figure 2: Graphical Model for tracking with ranked signals

To investigate the question of the practical advantage of our ordering constraint, which is a key difference from previous work, we consider a baseline model by dropping the ranking assumption. Thus, any permutation is allowed as a possible signal to signal source assignment, and is given the same weight. This approach leads to a mixture over possible assignments and we compare against this baseline in our experiments.

2 GRAPHICAL MODEL

Our proposed graphical model for tracking with ranked signals is shown in Figure 2. While Figure 2 expresses our model as a directed graphical model (also known as a Bayesian network), we will specify the probability distribution not in terms of conditional probabilities but in terms of local factors in the corresponding factor graph (see Figure 9 in appendix), since these will be used for inference in the sequel. For simplicity of description, we assume that target is in a 2-dimensional space, however it must be noted that our model applies to the general case when the target is in a d-dimensional space. We now describe each node and factor in Figure 2.

At each time step, the target records the first r signals emitted from s signal sources. It is possible that a signal from a particular signal source may be lost and not detectable at the target for the target to record. It is also possible that the target may record a spurious signal if the spurious signal is ranked higher than a true signal.

We indicate the position of the target by $X_t \in \mathbb{R}^2$. We model the system's dynamics as $X_{t+1} \sim \mathcal{N}(X_t, \mathbf{\Sigma})$, so that we have the corresponding factor $f(X_t, X_{t+1})$ given by

$$f(X_t, X_{t+1}) = (2\pi)^{-1}|\Sigma|^{-1/2} \exp\left(-\frac{1}{2}\Delta_t^T \Sigma^{-1} \Delta_t\right) \tag{1}$$

where $\Delta_t = X_{t+1} - X_t$.

We denote the signal characteristic (providing information about the target) from signal source i by $t_i \in \mathbb{R}, (1 \leq i \leq s)$. Note that in the context of RAFOS float tracking, this would correspond to the arrival time of the sound signal from beacon i. Some of these signals might be lost as noted in the introduction, so that we use Bernoulli random variables $m_i \in \{0, 1\}$ to indicate whether or not the signal from signal source i was lost. Thus, if $m_i = 1$, the signal was lost and t_i is set to ∞ (so that signal source i will be ignored in (3)). Otherwise, t_i follows some given distribution depending on signal source i and current the target location X_t. For example, in the RAFOS tracking problem, t_i follows a Gaussian distribution centered on the time taken for a sound signal to travel from location b_i to location X_t, given by $\frac{\|X_t - b_i\|_2}{v_s}$, where b_i is the beacon's location and v_s is the speed of sound. We capture this interaction between the state X_t, the estimated arrival time t_i, and the lost signal indicator m_i via the factors $g_i(X_t, t_i, m_i)$ $(1 \leq i \leq s)$:

$$g_i(X_t, t_i, m_i) = \begin{cases} \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{\left(t_i - \frac{\|X_t - b_i\|_2}{v_s}\right)^2}{2\sigma^2}} & \text{if } m_i = 0 \\ 1 & \text{if } m_i = 1 \end{cases} \tag{2}$$

Note that $t_i \in \mathbb{R}, (1 \leq i \leq s)$ are not directly observed, and hence latent variables in the model. Next, we consider the *observed* signals which we denote by $T_i (1 \leq i \leq r)$. Since these signals are stored sequentially, we have that $T_1 < T_2 < \cdots < T_r$ by our ranking assumption. But some of these signals may not correspond to signals from actual signal sources at all, and could be purely due to *clutter*, as noted in the introduction. We thus use Bernoulli random variables $c_i \in \{0, 1\}$ to indicate whether or not the signal T_i corresponds to clutter. By the assumptions made in the previous section, if the signal from the signal source that is supposed to be recorded first is not lost, i.e. $t_i < \infty$, and moreover the value stored as T_i is not junk, i.e. $c_i = 1$ then T_i *must* be the minimum of $\{t_j\}_{j=1}^{i}$. Otherwise, a junk clutter value is recorded. We assume these clutter values follow some fixed distribution, which we denote using the random variables n_i. For instance, in the RAFOS float tracking problem we assume junk clutter value are distributed *uniformly* over a specified interval.

We represent this interaction of T_j with the clutter, lost/missing, and signal variables, via the factor $f(T_j, c_1, c_2, \dots, c_j, n_j, t_1, m_1, t_2, m_2, \dots, t_s, m_s)$:

$$\begin{aligned} &f(T_j, c_1, c_2, \dots, c_j, n_j, t_1, m_1, t_2, m_2, \dots, t_s, m_s) \\ &= \begin{cases} \delta(T_j - t_{(j)}) & \text{if } c_j = 0 \\ \delta(T_j - n_j) & \text{if } c_j = 1 \end{cases} \end{aligned} \tag{3}$$

where we use $t_{(j)}$ to denote the $(j - \sum_{l=1}^{j-1} c_l)^{\text{th}}$ smallest element of the set $\{t_l\}$. This is a degenerate conditional probability distribution for T_i with mass only at the appropriate t_i or n_i. Recall that δ is the Dirac delta function, which has the following properties:

$$\delta(x) = \begin{cases} 0 & \text{if } x \neq 0 \\ \text{undefined} & \text{if } x = 0 \end{cases}, \int \delta(x)\mathrm{d}x = 1,$$
$$\int f(x)\delta(x)\mathrm{d}x = f(0) \tag{4}$$

2.1 INFERENCE WITH min FACTORS

The key novelty of our graphical model are the high order factors f_j ,which depend on the r first ranked signals. We consider this model in the simplified case, with no missing signals and no clutter, and show how inference via message passing can be performed in this model. In this setting, the factor $f_j(T_j, t_1, t_2, \dots, t_s)$ is given by

$$f_j(T_j, t_1, t_2, \dots, t_s) = \delta(T_j - t_{(j)}) \tag{5}$$

where $t_{(j)}$ is the j^{th} minimum element of $\{t_1, t_2, \dots, t_s\}$.

Direct computation of messages for high order factors in general requires computing an $s-1$-dimensional integral. However, our f_j, which correspond to the j-th minimum function, can be rewritten as a sum of products as,

$$f_j = \sum_{k=1}^{s} \delta(t_k - T_j) \sum_{(A,B)\in\mathcal{S}_k} \prod_{a\in A} \mathbf{1}(t_a < T_j) \prod_{b\in B} \mathbf{1}(t_b > T_j) \tag{6}$$

where $\mathcal{S}_k = \{(A,B) \subseteq [s]\times[s] : A \cup B = [s]\setminus\{k\}, A\cap B = \varnothing, |A| = j-1, |B| = s-j\}$ and $[s] = \{1, 2, \dots, s\}$

Then, as we show in the appendix, the multidimensional integral, comprising messages in a message passing inference algorithm, can be computed using only one-dimensional integrals and turn out to have the form:

$$\begin{aligned}\mu_{f_j\to t_i}(t_i) =& \delta(t_i - T_j)h_1(T_j) + \mathbf{1}(t_i < T_j)h_2(T_j) \\ &+ \mathbf{1}(t_i > T_j)h_3(T_j)\end{aligned} \tag{7}$$

where each h_j can be computed in $O(s^r)$ time via dynamic programming. When r is fixed, as in the RAFOS float tracking problem – since with $r = 3$ at most three known nearest beacons we can obtain the floats position and the overall number of beacons s is irrelevant – these messages can effectively be computed in polynomial time. This result is in the vein of previous work on reductions for high order potentials such as (Tarlow *et al.* , 2010). Handling high order min factors in this way is a novel result and can potentially be applied to other problems where such ranking factors are involved. Due to their technical complexity however, we defer further discussion of these factors and the details of the message passing algorithm to Appendix A.2. Instead, in the next section, we consider a reduction of our graphical model, with many such min factors, to an auxiliary model with a bipartite graph structure, which results in easier to implement algorithms.

2.2 REDUCTION TO BIPARTITE FACTORS

In this section, we show that our Bayesian network can be represented via a simpler conditional random field with factors that enforce a bipartite matching constraint. We do this by introducing latent variables S_i, R_j (Figure 3) to represent the association between signal sources and signals.

One large factor. First, consider the simplified case with no clutter and no missing signals. In this case, our graphical model is equivalent to one with the following large factor f connecting the $t_1, t_2, \dots, t_s$ and $T_1, T_2, \dots, T_r$ (instead of the f_j of (3) connecting t_j to the $\{T_i\}$), $f(t_1, t_2, \dots, t_s, T_1, T_2, \dots, T_r)$ as

$$\begin{aligned} &f(t_1, t_2, \dots, t_s, T_1, T_2, \dots, T_r) \\ &= \sum_{\substack{\Pi\in\{r\text{-permutations} \\ \text{of } \{1,2,\dots,s\}\}}} \left\{ \prod_{j=1}^{r} \delta(t_{\Pi(j)} - T_j) \prod_{k\notin\Pi} \mathbf{1}(t_k > T_r) \right\} \end{aligned} \tag{8}$$

We prove this by showing that the factor f defined in (8) above satisfies:

$$f = \prod_{j=1}^{r} f_j \tag{9}$$

where $f_1, f_2, \dots, f_j$ are the factors in eq (3). For any $\{t_1, t_2, \dots, t_s\}$, if the r smallest elements are not $T_1, T_2, \dots, T_r$ then we can see that $f = 0$ and $\prod_{j=1}^{r} f_j = 0$. When the r smallest elements of $\{t_1, t_2, \dots, t_s\}$ are $T_1, T_2, \dots, T_r$, where $t_{i_1} = T_1, t_{i_2} = T_2, \dots, t_{i_r} = T_r$, then we can see that $f = \prod_{j=1}^{r} \delta(t_{i_j} - T_j) = \prod_{j=1}^{r} f_j$.

Auxiliary Bipartite Graph. f has a special structure in that it sums over *partial* matchings and for each matching, it can be written as a product of pairwise factors, represented by the snippet of a pairwise conditional random field shown in Figure 3. To represent the matching, we use variables S_i ($1 \le i \le s$) corresponding to signal sources: these take values in $\{1, 2, \dots, r, \mathtt{L}\}$, indicating which of the r received signals S_i corresponds to, or a value of $\mathtt{L}$ if the signal was *late*, i.e. not in the first r and was thus not received. On the other side of the matching, we have variables R_j ($1 \le j \le r$) corresponding to received signals T_j: these take values in $\{1, 2, \dots, s\}$ indicating which of the signal sources T_j originated from. Thus, T_j is the signal corresponding to source R_j, while T_{S_i} is the signal corresponding to source i.

The factors $f(S_i, R_j)$ enforce these bipartite matching constraint:

$$f(S_i, R_j) = \begin{cases} 0 & \text{if } S_i = j, R_j \neq i \text{ or} \\ & \quad S_i \neq j, R_j = i \\ 1 & \text{otherwise} \end{cases} \tag{10}$$

And the factor $f(S_i)$ is given by

$$f(S_i) = \begin{cases} \delta(t_i - T_{S_i}) & \text{if } S_i \neq \mathtt{L} \\ \mathbf{1}(t_i > T_r) & \text{if } S_i = \mathtt{L} \end{cases} \tag{11}$$

Then, on marginalizing S_i and R_j, we obtain the factor f as in equation (8). Thus, we can replace the factor f of equation (8) by this auxiliary graphical model and perform message passing on this graphical model. This leads to a model similar to the pairwise models proposed to approximate the permanent, which shows up in the usual data association problem (Pasula *et al.* , 1999; Huang & Jebara, 2009; Oh *et al.* , 2009; Chertkov *et al.* , 2010; Vontobel, 2013). We can now consider the difference between our bipartite model and the one used to approximate the permanent (in (Huang & Jebara, 2009)): the sum in eq (8) is over partial permutations while that in (Huang & Jebara, 2009) involves full permutations. Our model consequently involves an additional value L to indicate that a signal is "late", leading to the node factors as shown in eq (11) which use indicator functions to indicate that a sent signal may not be in the top r ranked signals, and thus encodes the $\min$ factors required by our model.

Generalization to Factors with Missing Values, Clutter. We now generalize this argument to yield a similar conditional random field with pairwise factors for the complete model with missing values and clutter. Each S_i ($1 \leq i \leq s$) takes values $1, 2, \ldots, r, \mathtt{L}, \mathtt{M}$, where a number indicates which signal it corresponds to, L indicates that the signal is ranked too low to be observed, and M indicates that the signal is missing and undetectable at the receiver. Each R_j ($1 \leq j \leq r$) takes values $1, 2, \ldots, s, \mathtt{C}$, where a number indicates the signal's source, and C indicates it's clutter. The factors $f(S_i, R_j)$ still enforce the bipartite matching constraint of equation (10). The factors $f(n_j, R_j)$ are given by

$$f(n_j, R_j) = \begin{cases} P_C \delta(n_j - T_j) & \text{if } R_j = \mathtt{C} \\ 1 - P_C & \text{otherwise} \end{cases} \tag{12}$$

where P_N is the probability that the signal is clutter.

And the factor $f(t_i, S_i)$ is given by

$$f(t_i, S_i) = \begin{cases} 1 - P_D & \text{if } S_i = \mathtt{M} \\ P_D \mathbf{1}(t_i > T_r) & \text{if } S_i = \mathtt{L} \\ P_D \delta(t_i - T_{S_i}) & \text{otherwise} \end{cases} \tag{13}$$

where P_D is the probability that the signal can be detected

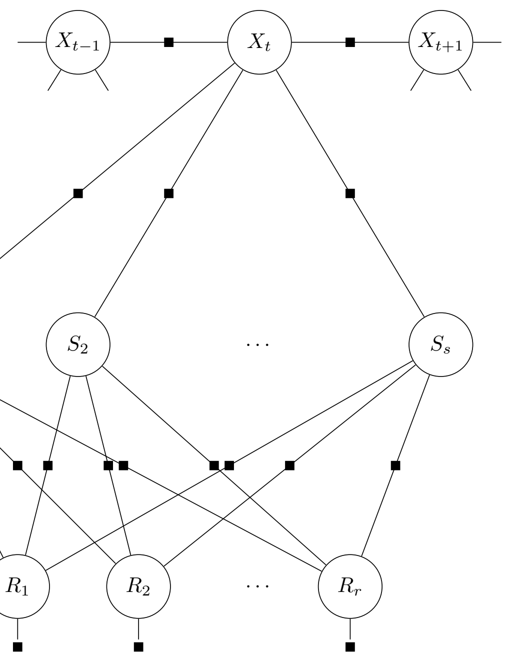

Figure 3: A conditional random field representing the distribution of r earliest arriving signals

Finally, we marginalize out $t_1, t_2, \ldots, t_s$ and $n_1, n_2, \ldots, n_r$, so that we can represent the conditional distribution as a simpler pairwise conditional random field as in Figure 3, where no factor uses the delta function, simplifying computation.

In the conditional random field shown in Figure 3, the factor $f(X_t, S_i)$ is given by

$$f(X_t, S_i) = \begin{cases} 1 - P_D & \text{if } S_i = \mathtt{M} \\ P_D \int_{T_r}^{+\infty} p_i(T|X_t)\, dT & \text{if } S_i = \mathtt{L} \\ P_D p_i(T_{S_i}|X_t) & \text{otherwise} \end{cases} \tag{14}$$

In the RAFOS tracking problem $p_i(T|X_t) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left\{ -\frac{\left(T - \frac{\|X_t - b_i\|_2}{v_s}\right)^2}{2\sigma^2} \right\}$ is the distribution of signal arrival time T given the location X_t of the target, and the factor $f(X_t, R_j)$ is given by

$$f(R_j) = \begin{cases} P_C p_N(T_j) & \text{if } R_j = \mathtt{C} \\ 1 - P_C & \text{otherwise} \end{cases} \tag{15}$$

where $p_N(T)$ is the distribution of clutter signal arrival time.

Alternative Derivation of the Bipartite Model Here we show an alternative and more intuitive derivation of the bipartite model assuming that the number of clutter signals follows a Poisson distribution with parameter λ and clutter signals are independent and identically distributed. We also assume that the distribution of signals from different signal sources are independent, and whether or not a signal is missing and non-detectable at the target is independent of all other signals.

Assume that the target's location is X, and there are N clutter signals, the likelihood of observing $T_1 \leq T_2 \leq \cdots \leq T_r$ with labels $S_1, S_r, \ldots, S_s$ and $R_1, R_2, \ldots, R_r$ is proportional to

$$\begin{aligned}\ell_N = & \frac{N!}{(N-\sum \mathbf{1}(R_j=\mathtt{C}))!}(1-P_D)^{\sum \mathbf{1}(S_i=\mathtt{M})} P_D^{\sum \mathbf{1}(S_i\neq\mathtt{M})} \\ & \left(\int_{T_r}^{+\infty} p_N(t)\,dt\right)^{N-\sum \mathbf{1}(R_j=\mathtt{C})} \prod_{R_j=\mathtt{C}} p_N(T_j) \\ & \prod_{S_i\neq \mathtt{L},\mathtt{M}} p_i(T_{S_i}|X) \prod_{S_i=\mathtt{L}} \int_{T_r}^{+\infty} p_i(t|X)\,dt \end{aligned} \tag{16}$$

Marginalizing out N we then have

$$\begin{aligned}\sum_N e^{-\lambda}\frac{\lambda^N}{N!}\ell_N \propto & \lambda^{\sum \mathbf{1}(R_j=\mathtt{C})}(1-P_D)^{\sum \mathbf{1}(S_i=\mathtt{M})} \\ & P_D^{\sum \mathbf{1}(S_i\neq\mathtt{M})} \prod_{R_j=\mathtt{C}} p_N(T_j) \prod_{S_i\neq \mathtt{L},\mathtt{M}} p_i(T_{S_i}|X) \\ & \prod_{S_i=\mathtt{L}} \int_{T_r}^{+\infty} p_i(t|X)\,dt \end{aligned} \tag{17}$$

Thus if we set

$$P_C = \frac{\lambda}{1+\lambda} \tag{18}$$

then it is clear that we can represent this distribution using the conditional random field shown in Figure 3.

2.3 INFERENCE

We use a hybrid particle filtering and belief propagation approach. For propagating information across time steps, i.e. between the nodes X_{t-1} and X_t, we use particle filtering (Doucet & Johansen, 2009). Particle based methods are simple to implement for tracking with range measurements (Cevher *et al.* , 2006). Such methods have also been widely used in robotics (Thrun *et al.* , 2005).

The distribution for each X_t is represented by a set of particles, we use message passing in the graphical model shown in Figure 3 to compute $P(T_1, T_2, \ldots, T_r|X_t)$, which is given by the partition function and can be approximated by the Bethe free energy (Huang & Jebara, 2009). In this graphical model, message passing equations have a simple form where each iteration is $O(sr)$ (Huang & Jebara, 2009; Williams & Lau, 2010), and message passing converges to a unique fixed point (Huang & Jebara, 2009; Williams & Lau, 2010; Vontobel, 2013). Previous work (Huang & Jebara, 2009; Williams & Lau, 2010; Chertkov *et al.* , 2010) has shown this strategy to be very effective.

Algorithm 1 Algorithm for tracking with ranked signals

function TRACKWITHRANKEDSIGNALS(X_0, $\{T_1^{(i)} \leq T_2^{(i)} \leq \cdots \leq T_r^{(i)}\}_{i=1}^n$)
  $N_p \leftarrow$ number of particles to use
  $p[N_p] \leftarrow$ each particle is initialized to X_0
  $w[N_p]$ ▷ each particle's weight
  for $i = 1;\ i \leq n;\ ++i$ **do**
    for $j = 0;\ j \neq N_p;\ ++j$ **do**
      $p[j] \leftarrow$ SAMPLE($P(X_t|X_{t-1} = p[j])$)
      $w[j] \leftarrow P(T_1^{(i)}, T_2^{(i)}, \ldots, T_r^{(i)}|p[j])$ ▷ computed using belief propagation in the conditional random field in Figure 3
    end for
    $p \leftarrow$RESAMPLE(p,w) ▷ particle filter resampling
  end for
end function

An implementation of the tracking algorithm to estimate $X_1, X_2, \ldots, X_n$ where we are given the initial position of the target X_0, and n observations of ranked signals $\{T_1^{(i)} \leq T_2^{(i)} \leq \cdots \leq T_r^{(i)}\}_{i=1}^n$ is given in Algorithm 1.

3 EXPERIMENTS

We present experiments on two kinds of datasets: simulated data generated using our model (Figure 3) and real world data from RAFOS floats.

3.1 SIMULATIONS

We simulate tracking RAFOS floats using ranked continuous range measurements with our method. At each time step, the target records the arrival times of the first 4 signals emitted from 10 fixed beacons. There is a fixed probability that a signal may not be detectable at the target, and the target may record the arrival time of a spurious signal (clutter) if the spurious signal arrives before an actual signal.

3.1.1 SIMULATION OF TRACKING DIFFERENT TRAJECTORIES

Figure 4a shows our algorithm tracking a target moving in a straight line: the x and y axes correspond to the x and y coordinates of the 2D location of the signal. The corresponding signal arrival times are shown in Figure 4b. Similarly, Figure 5a shows our algorithm tracking a target moving in a spiral, signal arrival times are shown in Figure 5b. In the simulations, there are 10 signal sources and 4 signal arrival times. Clutter arrival times are uniformly distributed on a predefined interval, and the number of clutter signals follows a Poisson distribution with parameter 4, so that $P_C = \frac{4}{1+4} = 0.8$. The signal arrival time distribution p_i has parameter $\sigma = 0.02$, and $P_D = 0.7$. We can see that our algorithm correctly tracks the target when the

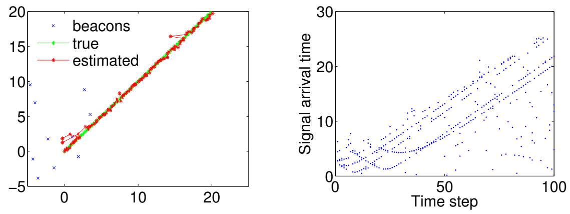

(a) Our algorithm's result (b) Signal arrival times

Figure 4: Our algorithm tracking a target moving in a straight line

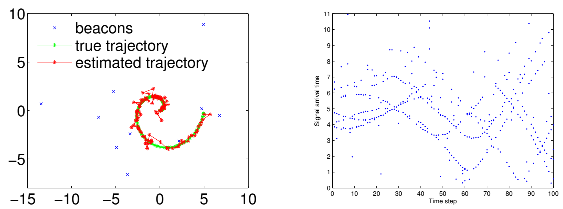

(a) Our algorithm's result (b) Signal arrival times

Figure 5: Our algorithm tracking a target moving in a spiral

association between signal arrival times and signal sources changes.

3.1.2 COMPARISON WITH THE BASELINE ALGORITHM

We compare our algorithm against a baseline algorithm which does not take into account that signal arrival times are ranked. The overall structure of the baseline model is the same as that of the graphical model shown in Figure 3, however the factor $f(X_t, S_i)$ is given by

$$f(X_t, S_i) = \begin{cases} P_D p_i(T_{S_i}|X_t) & \text{if } S_i \neq \mathtt{L} \text{ and } S_i \neq \mathtt{M} \\ 1 - P_D & \text{if } S_i = \mathtt{L} \text{ or } S_i = \mathtt{M} \end{cases} \tag{19}$$

so that the implicit ranking of signal arrival times is not taken into account.

Figure 6b shows the mean squared error against the noise level when tracking the straight trajectory shown in Figure 6a. The noise level is indicated by a parameter λ such that $P_C = \frac{\lambda}{1+\lambda}$. The parameter λ corresponds to the Poisson parameter for the number of clutter signals (see Appendix). Signal arrival time distribution p_i has parameter $\sigma = 0.02$, and $P_D = 0.7$. We can see that our algorithm outperforms the baseline algorithm.

3.2 RAFOS FLOAT DATA

We present tracking results on real world RAFOS float data which consists of ranked continuous signal arrival times. We compare the results of our algorithm against hand labeled data where each signal is hand labeled with a particular signal source or as clutter. In RAFOS float missions, after a float is released at a known location to start its mission, it records a fixed number of arrival times from a set of beacons whose positions are fixed and known. The dataset presented here are collected in the DIMES (Diapycnal and Isopycnal Mixing Experiment in the Southern Ocean) project (Hancock & Speer, 2013), which is aimed at measuring diapycnal and isopycnal mixing in the Southern Ocean, along the tilting isopycnals of the Antarctic Circumpolar Current. In our dataset, the float records 4 earliest signal arrival times, and there are 10 beacons. The floats are first released off the western coast of South America. There were just two floats — #767 and #811 — for which at least three beacon signals were recorded at all time steps (note that at least three beacon signals are required to uniquely identify the 2D position of the float), and accordingly, we present our results on tracking these two floats. Their signal arrival times are shown in Figure 1. We compare our algorithm's results to a manual tracking procedure where each signal is hand labeled to its source or as noise based on intuition and some apriori knowledge of the physical factors (geometry of ocean basin and acoustic array, basic knowledge of current directions etc.) at play.

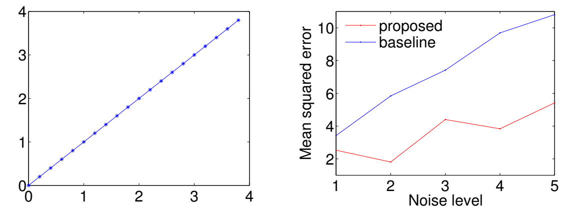

(a) A straight trajectory used to compare our algorithm and the baseline algorithm (b) Mean squared error of our algorithm and the baseline algorithm for different noise levels λ

Figure 6: Simulation comparing our proposed algorithm and the baseline

Note that the *true* locations of the RAFOS floats are not known and that the hand labeled data only contains the associations between signal arrival times and their corresponding signal sources (or whether they are clutter). Hence, we obtain a trajectory for the hand labeled data using a simple particle filter. At each time step $T_1 \leq T_2, \leq \cdots \leq T_r$ and $R_1, R_2, \ldots, R_r$ are known, so the weight of a particle at X is given by

$$w(X) \propto \prod_{1 \leq j \leq r} p_{R_j}(T_j|X) \tag{20}$$

When running our proposed algorithm and the one using hand labeled data for RAFOS float tracking, we set $P_C = 0.5$, $P_D = 0.9$, $\sigma = 5$. Although in the real environment the speed of sound depends on the depth of the RAFOS float and other environmental factors, in our experiments we set $v_s = 1.5$ km/s across all cases.

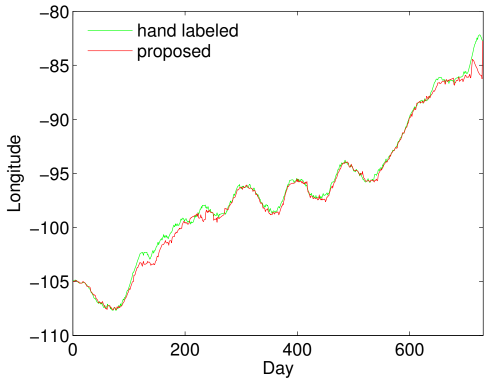

(a) Longitude

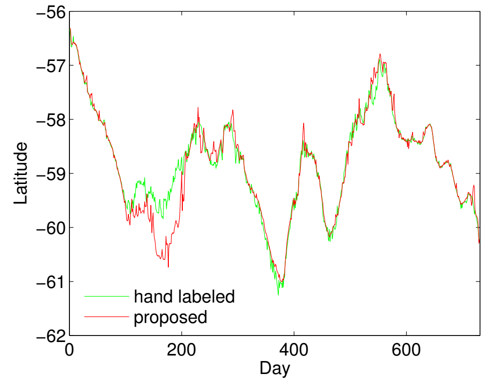

(b) Latitude

Figure 7: Trajectory of float #767 estimated by our proposed algorithm versus using hand labeled data

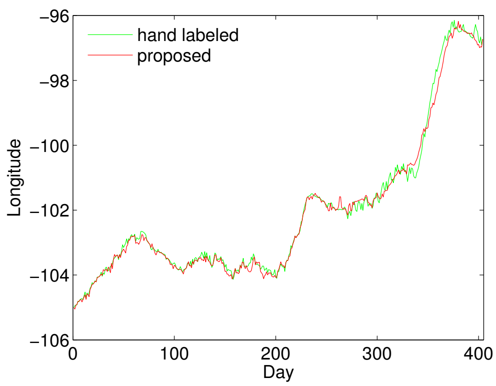

(a) Longitude

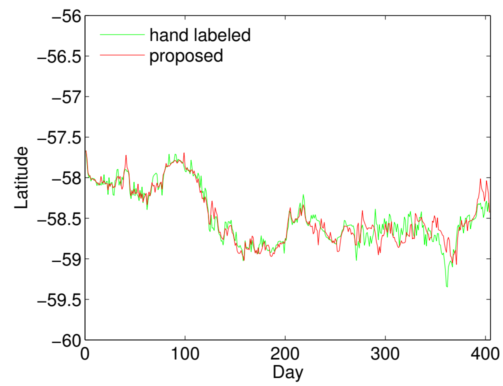

(b) Latitude

Figure 8: Trajectory of float #811 estimated by our proposed algorithm versus using hand labeled data

The results for longitude and latitude of float #767 and #811 estimated by our algorithm versus using hand labeled data are shown in Figure 7 and Figure 8 respectively. Additional results for other floats in the DIMES project are presented in the appendix. We observe good agreement between the tracks estimated using hand labeled data and using our method indicating that our method recovers the associations. We can also see that our algorithm continues to track the target when the set of associated signal sources changes.

4 CONCLUSIONS

We have presented a novel graphical model approach for the problem of tracking with ranked signals which was able to capture certain problem specific features easily. The key novelty in the model, from a machine learning point of view, was the presence of ranking-based factors. We have provided a novel bipartite construction which allows for easy inference via message passing for such factors. While our model and approach were motivated by a particular application, the contributions of this paper should be applicable to other problems where ranking based high-order factors are involved.

We experimentally showed that our method effectively leverages the varied problem constraints to improve tracking accuracy over baseline tracking methods. We applied our method for tracking with ranked signals to a key Oceanography problem of tracking RAFOS floats in the ocean and thus provide an automated solution to a problem which has previously required significant manual effort.

ACKNOWLEDGEMENTS

T.L., H.P., and P.R. acknowledge the support of ARO via W911NF-12-1-0390 and NSF via IIS-1149803, IIS-1320894, IIS-1447574, and DMS-1264033. K.S. and D.B. acknowledge support from NSF OCE 1231803.

References

Bar-Shalom, Y., Daum, F., & Huang, J. 2009. The probabilistic data association filter. *Control Systems, IEEE*, **29**(6), 82–100.

Bar-Shalom, Yaakov. 1987. *Tracking and data association*. Academic Press Professional, Inc.

Blackman, Samuel S. 2004. Multiple hypothesis tracking for multiple target tracking. *Aerospace and Electronic Systems Magazine, IEEE*, **19**(1), 5–18.

Cevher, Volkan, Velmurugan, Rajbabu, & McClellan, James H. 2006. A range-only multiple target particle filter tracker. *Pages IV–IV of: Acoustics, Speech and Signal Processing, 2006. ICASSP 2006 Proceedings. 2006 IEEE International Conference on*, vol. 4. IEEE.

Chertkov, M, Kroc, L, Krzakala, F, Vergassola, M, & Zdeborová, L. 2010. Inference in particle tracking experiments by passing messages between images. *Proceedings of the National Academy of Sciences*, **107**(17), 7663–7668.

Doucet, Arnaud, & Johansen, Adam M. 2009. A tutorial on particle filtering and smoothing: Fifteen years later. *Handbook of Nonlinear Filtering*, **12**, 656–704.

Hancock, C., & Speer, K. 2013. Diapycnal and Isopycnal Mixing Experiment in the Southern Ocean, RAFOS Float Data Report, Marine Field Group, FSU, January 2013.

Huang, Bert, & Jebara, Tony. 2009. Approximating the permanent with belief propagation. *arXiv preprint arXiv:0908.1769*.

Huang, Jonathan, Guestrin, Carlos, & Guibas, Leonidas. 2009. Fourier theoretic probabilistic inference over permutations. *The Journal of Machine Learning Research*, **10**, 997–1070.

Jerrum, Mark, Sinclair, Alistair, & Vigoda, Eric. 2004. A polynomial-time approximation algorithm for the permanent of a matrix with nonnegative entries. *Journal of the ACM (JACM)*, **51**(4), 671–697.

Kondor, Risi, Howard, Andrew, & Jebara, Tony. 2007. Multi-object tracking with representations of the symmetric group. *Pages 211–218 of: International Conference on Artificial Intelligence and Statistics*.

Oh, Songhwai, Russell, Stuart, & Sastry, Shankar. 2009. Markov chain Monte Carlo data association for multi-target tracking. *Automatic Control, IEEE Transactions on*, **54**(3), 481–497.

Pasula, Hanna, Russell, Stuart, Ostland, Michael, & Ritov, Yaacov. 1999. Tracking many objects with many sensors. *Pages 1160–1171 of: IJCAI*, vol. 99.

Rossby, T, Dorson, D, & Fontaine, J. 1986. The RAFOS system. *Journal of Atmospheric and Oceanic Technology*, **3**(4), 672–679.

Tarlow, Daniel, Givoni, Inmar E, & Zemel, Richard S. 2010. Hop-map: Efficient message passing with high order potentials. *Pages 812–819 of: International Conference on Artificial Intelligence and Statistics*.

Thrun, Sebastian, Burgard, Wolfram, & Fox, Dieter. 2005. *Probabilistic robotics*. MIT press.

Valiant, Leslie G. 1979. The complexity of computing the permanent. *Theoretical computer science*, **8**(2), 189–201.

Vontobel, Pascal O. 2013. The Bethe permanent of a non-negative matrix. *Information Theory, IEEE Transactions on*, **59**(3), 1866–1901.

Williams, Jason L, & Lau, Roslyn A. 2010. Convergence of loopy belief propagation for data association. *Pages 175–180 of: Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP), 2010 Sixth International Conference on*. IEEE.

Wooding, Christine M, Furey, Heather H, & Pacheco, Marguerite A. 2005. *RAFOS float processing at the Woods Hole Oceanographic Institution*. Tech. rept. Woods Hole Oceanographic Institution.

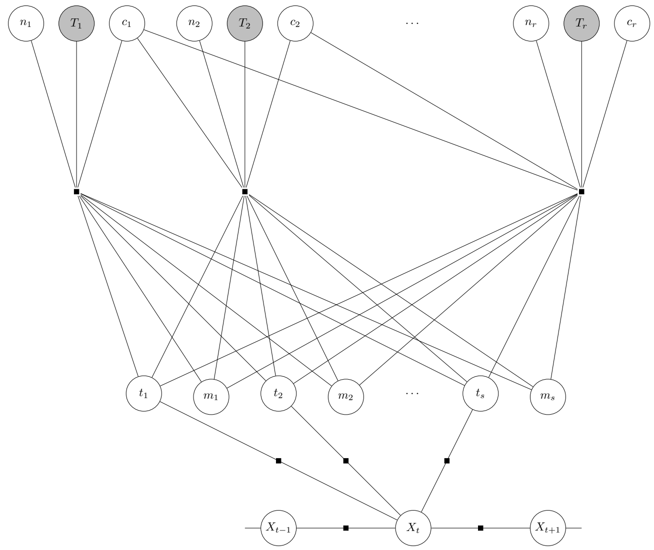

Figure 9: Factor graph corresponding to Figure 2

A APPENDIX

A.1 MESSAGE PASSING EQUATIONS FOR BIPARTITE GRAPH

We use message passing for inference in the factor graph shown in Figure 10 where at time step t $T_1, T_2, \ldots, T_r$ are observed. In the following description of message passing at time step t we sometimes omit the t subscript for notational conveniences.

For each S_i, the message $\nu_{S_i \to S_i,R_j}$ to the factor $f(S_i, R_j)$ is

$$\nu_{S_i \to S_i,R_j} = \mu_{S_i,X_t,T_r \to S_i} \prod_{\substack{1 \le J \le r \\ J \neq j}} \mu_{S_i,R_J \to S_i} \tag{21}$$

the message $\mu_{S_i,R_j \to S_i}$ from the factor $f(S_i, R_j)$ is

$$\mu_{S_i,R_j \to S_i} = \sum_{R_j} f(S_i, R_j)\, \nu_{R_j \to S_i,R_j} \tag{22}$$

the message $\nu_{S_i \to S_i,X_t,T_r}$ to the factor $f(S_i, X_t, T_r)$ is

$$\nu_{S_i \to S_i,X_t,T_r} = \prod_{1 \le J \le r} \mu_{S_i,R_J \to S_i} \tag{23}$$

the message $\mu_{S_i,X_t,T_r \to S_i}$ from the factor $f(S_i, X_t, T_r)$ is

$$\mu_{S_i,X_t,T_r \to S_i} = \tag{24}$$

$$\int f(S_i, X_t, T_r)\, \nu_{X_t \to S_i,X_t,T_r} \nu_{T_r \to S_i,X_t,T_r}\, dX_t \tag{25}$$

For each R_j, the message $\nu_{R_j \to S_i,R_j}$ to the factor

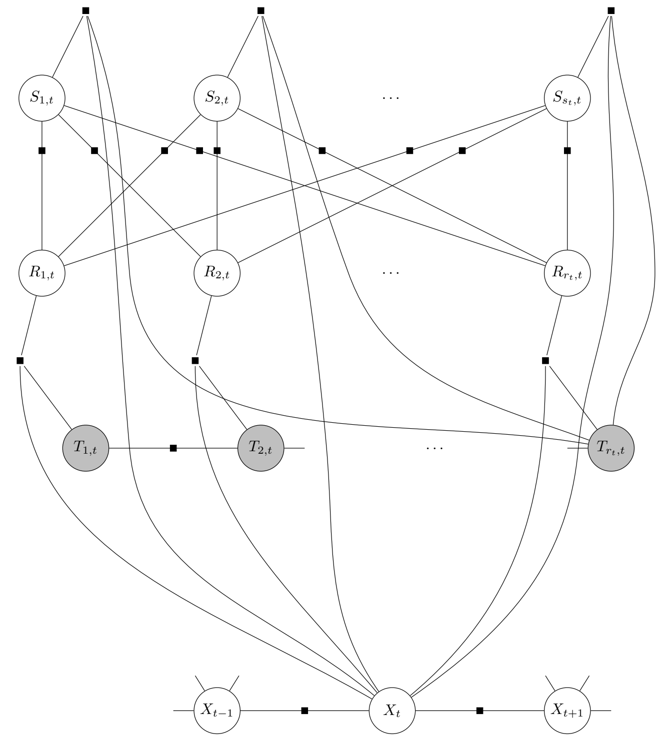

Figure 10: A factor graph representing the distribution of r earliest arriving signals in the bipartite model

$f(S_i, R_j)$ is

$$\nu_{R_j \to S_i,R_j} = \tag{26}$$

$$\mu_{R_1,R_2,\ldots,R_r,N \to R_j} \mu_{T_j,R_j,X_t \to R_j} \prod_{\substack{1 \le I \le s \\ I \neq i}} \mu_{S_I,R_j \to R_j} \tag{27}$$

the message $\mu_{S_i,R_j \to R_j}$ from the factor $f(S_i, R_j)$ is

$$\mu_{S_i,R_j \to R_j} = \sum_{S_i} f(S_i, R_j)\, \nu_{S_i \to S_i,R_j} \tag{28}$$

the message $\nu_{R_j \to T_j,R_j,X_t}$ to the factor $f(T_j, R_j, X_t)$ is

$$\nu_{R_j \to T_j,R_j,X_t} = \mu_{R_1,R_2,\ldots,R_r,N \to R_j} \prod_{1 \le I \le s} \mu_{S_I,R_j \to R_j} \tag{29}$$

the message $\mu_{T_j,R_j,X_t \to R_j}$ from the factor $f(T_j, R_j, X_t)$ is

$$\mu_{T_j,R_j,X_t \to R_j} = \int f(T_j, R_j, X_t)\, \nu_{X_t \to T_j,R_j,X_t}\, dX_t \tag{30}$$

the message $\nu_{R_j \to R_1,R_2,\ldots,R_r,N}$ to the factor $f(R_1, R_2, \ldots, R_r, N)$ is

$$\nu_{R_j \to R_1,R_2,\ldots,R_r,N} = \mu_{T_j,R_j,X_t \to R_j} \prod_{1 \le I \le s} \mu_{S_I,R_j \to R_j} \tag{31}$$

the message $\mu_{R_1,R_2,\ldots,R_r,N\to R_j}$ from the factor $f(R_1,R_2,\ldots,R_r,N)$ is

$$\mu_{R_1,R_2,\ldots,R_r,N\to R_j} = \tag{32}$$

$$\sum_{\{R_1,\ldots,R_r\}\setminus\{R_j\}} \sum_N f(R_1,\ldots,R_r,N) \tag{33}$$

$$\nu_{N\to R_1,\ldots,R_r,N} \prod_{\substack{1\le J\le r\\ J\neq j}} \nu_{R_J\to R_1,\ldots,R_r,N} \tag{34}$$

For N, the message $\nu_{N\to N}$ to the factor N is

$$\nu_{N\to N} = \mu_{R_1,R_2,\ldots,R_r,N\to N}\mu_{N,T_r\to N} \tag{35}$$

the message $\mu_{N\to N}$ from the factor N is

$$\mu_{N\to N} = \sum_N f(N) \tag{36}$$

the message $\nu_{N\to R_1,R_2,\ldots,R_r,N}$ to the factor $f(R_1,R_2,\ldots,R_r,N)$ is

$$\nu_{N\to R_1,R_2,\ldots,R_r,N} = \mu_{N\to N}\mu_{N,T_r\to N} \tag{37}$$

the message $\mu_{R_1,R_2,\ldots,R_r,N\to N}$ from the factor $f(R_1,R_2,\ldots,R_r,N)$ is

$$\mu_{R_1,R_2,\ldots,R_r,N\to N} = \tag{38}$$

$$\sum_{R_1,\ldots,R_r} f(R_1,\ldots,R_r,N) \prod_{1\le J\le r} \nu_{R_J\to R_1,\ldots,R_r,N} \tag{39}$$

the message $\nu_{N\to N,T_r}$ to the factor $f(N,T_r)$ is

$$\nu_{N\to N,T_r} = \mu_{N\to N}\mu_{R_1,R_2,\ldots,R_r,N\to N} \tag{40}$$

the message $\mu_{N,T_r\to N}$ from the factor $f(N,T_r)$ is

$$\mu_{N,T_r\to N} = \sum_N f(N,T_r) \tag{41}$$

For X_t, the message $\nu_{X_t\to X_{t-1},X_t}$ to the factor $f(X_{t-1},X_t)$ is

$$\nu_{X_t\to X_{t-1},X_t} = \tag{42}$$

$$\mu_{X_t,X_{t+1}\to X_t} \prod_{1\le I\le s} \mu_{S_I,X_t,T_r\to X_t} \prod_{1\le J\le r} \mu_{T_J,R_J,X_t\to X_t} \tag{43}$$

the message $\mu_{X_{t-1},X_t\to X_t}$ from the factor $f(X_{t-1},X_t)$ is

$$\mu_{X_{t-1},X_t\to X_t} = \int f(X_{t-1},X_t)\,\nu_{X_{t-1}\to X_{t-1},X_t}\; dX_{t-1} \tag{44}$$

the message $\nu_{X_t\to X_t,X_{t+1}}$ to the factor $f(X_t,X_{t+1})$ is

$$\nu_{X_t\to X_t,X_{t+1}} = \tag{45}$$

$$\mu_{X_{t-1},X_t\to X_t} \tag{46}$$

$$\prod_{1\le I\le s} \mu_{S_I,X_t,T_r\to X_t} \prod_{1\le J\le r} \mu_{T_J,R_J,X_t\to X_t} \tag{47}$$

the message $\mu_{X_t\to X_t,X_{t+1}}$ from the factor $f(X_t,X_{t+1})$ is

$$\mu_{X_t\to X_t,X_{t+1}} = \int f(X_t,X_{t+1})\,\nu_{X_{t+1}\to X_t,X_{t+1}}\; dX_{t+1} \tag{48}$$

the message $\nu_{X_t\to S_i,X_t,T_r}$ to the factor $f(S_i,X_t,T_r)$ is

$$\nu_{X_t\to S_i,X_t,T_r} = \tag{49}$$

$$\mu_{X_{t-1},X_t\to X_t}\mu_{X_t,X_{t+1}\to X_t} \tag{50}$$

$$\prod_{\substack{1\le I\le s\\ I\neq i}} \mu_{S_I,X_t,T_r\to X_t} \prod_{1\le J\le r} \mu_{T_J,R_J,X_t\to X_t} \tag{51}$$

the message $\mu_{S_i,X_t,T_r\to X_t}$ from the factor $f(S_i,X_t,T_r)$ is

$$\mu_{S_i,X_t,T_r\to X_t} = \sum_{S_i} f(S_i,X_t,T_r)\,\nu_{S_i\to S_i,X_t,T_r} \tag{52}$$

the message $\nu_{X_t\to T_j,R_j,X_t}$ to the factor $f(T_j,R_j,X_t)$ is

$$\nu_{X_t\to T_j,R_j,X_t} = \tag{53}$$

$$\mu_{X_{t-1},X_t\to X_t}\mu_{X_t,X_{t+1}\to X_t} \tag{54}$$

$$\prod_{1\le I\le s} \mu_{S_I,X_t,T_r\to X_t} \prod_{\substack{1\le J\le r\\ J\neq j}} \mu_{T_J,R_J,X_t\to X_t} \tag{55}$$

the message $\mu_{T_j,R_j,X_t\to X_t}$ from the factor $f(T_j,R_j,X_t)$ is

$$\mu_{T_j,R_j,X_t\to X_t} = \sum_{R_j} f(T_j,R_j,X_t)\,\nu_{R_j\to T_j,R_j,X_t} \tag{56}$$

A.2 MESSAGE PASSING FOR HIGH ORDER min FACTORS

Recall that the factor $f_j(T_j,t_1,t_2,\ldots,t_s)$ is given by

$$f_j(T_j,t_1,t_2,\ldots,t_s) = \delta(T_r - t_k) \tag{57}$$

where t_k is the j^{th} minimum element of $\{t_1,t_2,\ldots,t_s\}$. We denote this factor by f_j.

Direct computation of messages in this high order factor graph would require computing an $s-1$-dimensional integral. However, our f_j, which correspond to the j-th minimum function, can be rewritten as a sum of products as,

$$f_j = \sum_{k=1}^{s} \delta(t_k - T_j) \sum_{(A,B)\in\mathcal{S}_k} \prod_{a\in A} \mathbf{1}(t_a < T_j) \prod_{b\in B} \mathbf{1}(t_b > T_j) \tag{58}$$

where $\mathcal{S}_k = \{(A,B)\subseteq[s]\times[s] : A\cup B = [s]\setminus\{k\}, A\cap B = \varnothing, |A| = j-1, |B| = s-j\}$ and $[s] = \{1,2,\ldots,s\}$ The outer sum represents the s different cases where each element of $\{t_1, t_2, \ldots, t_s\}$ can be the j^{th} smallest. Suppose t_k is the j^{th} smallest and is equal to T_j. Then, the remaining $\{t_l | l \neq k\}$ are partitioned into 2 sets, where every t_l in one set is smaller than t_k and while each t_l in the other is larger. There are $\binom{s-1}{j-1}$ such partitions. Thus the f_j corresponds to a sum of products of $O(s\binom{s-1}{j-1})$ terms.

The message $\mu_{f_j \to t_i}(t_i)$ from the factor $f_j (1 \leq j \leq r)$ to the variable $t_i (1 \leq i \leq s)$ is given by:

$$\mu_{f_j \to t_i}(t_i) = \int \left(\prod_{\substack{1 \leq l \leq s \\ l \neq i}} \nu_{t_l \to f_j}(t_l) \right) f(T_j, t_1, t_2, \ldots, t_s) \underbrace{d \ldots t}_{\text{except } dt_i} \tag{59}$$

$$= \int \left(\prod_{\substack{1 \leq l \leq s \\ l \neq i}} \nu_{t_l \to f_j}(t_l) \right) \delta(T_j - t_k) \underbrace{d \ldots t}_{\text{except } dt_i} \tag{60}$$

where t_k is the j^{th} smallest element of $\{t_1, t_2, \ldots, t_s\}$, and $\nu_{t_l \to f_j}(t_l)$ is the message from t_l to f_j.

For computing $\mu_{f_j \to t_i}(t_i)$, f_j can be written as the sum of the following terms:

$$f_j = \delta(t_i - T_j) \sum_{A,B} \prod_{a \in A} \mathbf{1}(t_a < T_j) \prod_{b \in B} \mathbf{1}(t_b > T_j) \tag{61}$$

$$+ \sum_{k \neq i} \delta(t_k - T_j) \sum_{A,B} \prod_{a \in A} \mathbf{1}(t_a < T_j) \prod_{b \in B} \mathbf{1}(t_b > T_j) \tag{62}$$

Then, the multidimensional integral can be written as sum of products of unidimensional integrals. The final computation of the message requires a sum of $O(s\binom{s-1}{j-1})$ terms as,

$$\mu_{f_j \to t_i}(t_i) = \delta(t_i - T_j) h_1(T_j) + \mathbf{1}(t_i < T_j) h_2(T_j) + \mathbf{1}(t_i > T_j) h_3(T_j) \tag{63}$$

where

$$h_1(T_j) = \sum_{A,B} \prod_{a \in A} \left(\int_{-\infty}^{T_j} \nu_{t_a \to f_j}(t_a)\, dt_a \right) \tag{64}$$

$$\prod_{b \in B} \left(\int_{T_j}^{+\infty} \nu_{t_b \to f_j}(t_b)\, dt_b \right) \tag{65}$$

$$h_2(T_j) = \sum_{A,B,i \in A} \prod_{a \in A, a \neq i} \left(\int_{-\infty}^{T_j} \nu_{t_a \to f_j}(t_a)\, dt_a \right) \tag{66}$$

$$\prod_{b \in B} \left(\int_{T_j}^{+\infty} \nu_{t_b \to f_j}(t_b)\, dt_b \right) \tag{67}$$

$$h_3(T_j) = \sum_{A,B,i \in B} \prod_{a \in A} \left(\int_{-\infty}^{T_j} \nu_{t_a \to f_j}(t_a)\, dt_a \right) \tag{68}$$

$$\prod_{b \in B, b \neq i} \left(\int_{T_j}^{+\infty} \nu_{t_b \to f_j}(t_b)\, dt_b \right) \tag{69}$$

For r such factors f_j, if messages are computed directly, each iteration of message passing will require $O(\sum_{j=1}^{r} \binom{s}{j})$ computation. Note that only $2s$ unidimensional integrals need to be computed, and the remainder of the computation corresponds to computing the value of elementary symmetric polynomials, which corresponds to sums of all combinations. To compute a symmetric polynomial $\sum_{\substack{A \in \{1,2,\ldots,n\} \\ |A| = k}} \prod_{a \in A} c_a$ which sums over all k-combinations of $\{c_1, c_2, \ldots, c_n\}$, we can use dynamic programming to find the coefficient of x^k in $\prod_{i=1}^{n}(x + c_i)$, and this can be done in $O(n^2)$ time.

A.2.1 FULL MODEL WITH CLUTTER

We handle two kinds of systematic noise in this model: losses from the sender and clutter. Losses are handled by $m_1, m_2, \ldots, m_s$ in Figure 9.

Clutter can be incorporated in this model through the factor $f'_k(T_k, t_1, t_2, \ldots, t_s, J_1, J_2, \ldots, J_k)$ as

$$f'_k(.) = \begin{cases} \delta(T_r - t_l) & \text{if } J_k = 0 \\ 1 & \text{if } J_k = 1 \end{cases} \tag{70}$$

where t_l is the $(k - \sum_i J_i)^{\text{th}}$ minimum element of $\{t_1, t_2, \ldots, t_s\}$. This is identical to f from the previous section if the J_i are all zero. If $J_k = 1$, i.e. the current message is clutter, then we assume a uniform distribution over T_k. If some previous received message was clutter, T_k will take a lower minimum value.

Then, the factor can be written down in terms of the factors f, from the previous section, as $f'_k(T_k, t_1, t_2, \dots, t_s, J_1, J_2, \dots, J_k)$:

$$f'_k(.) = f_{k-\sum_i J_i}(T_k, t_1, t_2, \dots, t_s) \tag{71}$$

Then, the messages from f_k to t_i can be written as:

$$\nu'_{f'_k \to t_i} = \sum_{J_i} \nu_{f_{k-\sum_i J_i} \to t_i} \pi_l \mu_{J_l \to f'_k} \tag{72}$$

where the summation is over the values 0 or 1 for each J_i.

Messages from f'_k to J_l can be written as:

$$\nu'_{f'_k \to t_i} = \sum_{J_i, i \neq l} \int f_{k-\sum_i J_i} \mathrm{d}t_1 \dots \mathrm{d}t_s \tag{73}$$

Thus, we can precompute the messages for f_k in polynomial time, and we can compute these messages in $O(2^r)$ additional time.

A.3 ADDITIONAL EXPERIMENTAL RESULTS FOR RAFOS FLOAT DATA

Here we present more additional experimental results for tracking RAFOS floats using our proposed method. When there are at least three actual signal arrival times at each time step, such as float #767 and float #811 (Figure 1), it is possible to estimate a unique track for the float over the entire period of the float's mission (Figure 7 and 8). However, if at some point during a float's mission that there are only two actual signal arrival times for a certain period, then neither using hand labeled data nor our proposed method can uniquely determine the float's location.

An example for float #759 is given here. The signal arrival times for float #759 are shown in Figure 11, where there exists periods of time during float #759's mission when only at most two signal arrivals are available. As shown in Figure 12, we get different results in different runs of the simple particle filter algorithm using hand labeled data (blue), and our proposed algorithm agrees with hand labeled data when there are at least three signal arrival times available.

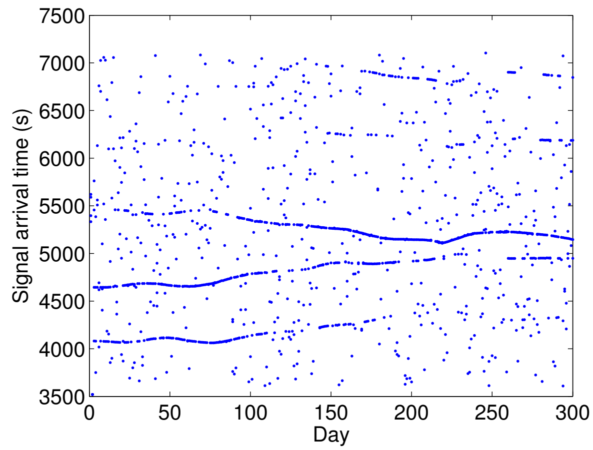

Figure 11: Observed signal arrival times for float #759 over the entire tracking period

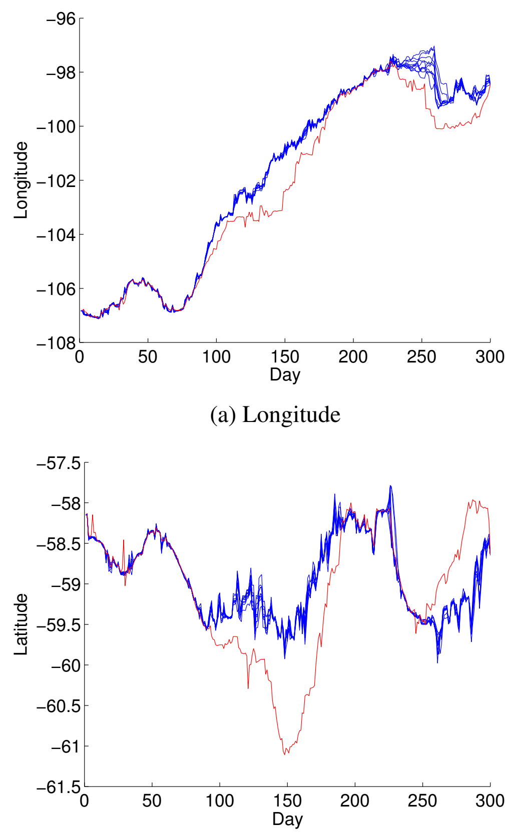

(a) Longitude

(b) Latitude

Figure 12: Results of different runs of the simple particle filter algorithm using hand labeled data (blue) versus our proposed algorithm (red) for float #759