Bayesian Track Correlation and Numbering

Egils Sviestins
Saab AB
SE-175 88 Järfälla
Sweden
egils.sviestins@saabtech.se
www.saabtech.se

Abstract – *In decentralized tracking a single situation picture is formed by correlating (associating) and merging tracks from different track sources. A problem rarely addressed in this connection is that of maintaining correct track numbers over time. The track numbers may have to be reassigned, for example when the track sources make mistakes like swapping tracks or picking up false measurements. This paper presents a coherent Bayesian approach to handling the dynamics in track correlation, where the basic idea is to consider track-to-target correlation instead of the more conventional track-to-track correlation.*

Keywords: track-to-track, correlation, association, Bayesian, numbering, distributed tracking, decentralized tracking.

1. Introduction

This paper deals with the art of forming a situation picture using data from several track sources, that is, decentralized tracking. Although decentralized tracking may not be the optimal architecture from the performance point of view, often this is the only possibility given sensors with built-in trackers or when there are severe restrictions on communications capacity. The challenge is thus to correlate (or associate) local tracks to local tracks, to determine which of them represent the same target, and then somehow form a single system track for each target. The track-to-track correlation problem has been studied a lot; for an overview, refer to e.g. [1].

But there is more to it. Even though the track-to-track correlation is or appears correct at one time, it does not necessarily mean it remains so in the future. Various phenomena do occur, apart from mistakes due to statistical fluctuations. Some examples are:

- Local tracks may be swapped at one source, and unswapped at another
- False local tracks may be correlated to genuine
- Tracks may split (more targets than originally detected)
- Tracks may diverge because of false measurement-to-track associations, or gaps in coverage
- Incorrectly diverged tracks may come close again

In practical applications it is very important that the track number given to the system track remains the same as long as the system track represents the same target, as usually all information about the target is connected to the track number. However, the work on track correlation rarely even mentions track numbering (for an exception, see [2]). Track numbering is usually left to ad hoc schemes.

The basic idea in this paper is that we should not consider track-to-track correlation in the first place. Instead track-to-target correlation is the fundamental issue, and here the target is represented by a target number. In case local tracks from different sources are assigned to the same target, it *follows* that the tracks are correlated to each other.

The advantage of this approach is that it makes it easy to formulate the various correlation hypotheses and to evaluate them recursively in a Bayesian manner, and also to take into account a lot of factors that are important for correct correlation.

The idea of correlating tracks to targets, rather than tracks to tracks has also appeared in [3], albeit in an entirely different context, namely Over-the-Horizon radar tracking.

2. Correlation Hypotheses

2.1 Track numbers vs. target numbers

Let S denote the set of sources contributing with local tracks (each source may or may not correspond to a sensor depending on whether the tracker is a single or multisensor tracker). Each source $s \in S$ contributes with a set L_s of local tracks. (In reality one is likely to consider subclusters of nearby tracks as input to the algorithm, and in that case S and the L_s refer only to one subcluster at a time.) Each element $L_{si} \in L_s$ is thus a local track, and we do not distinguish between the track and its track number. Suppose finally that the targets involved form a set T with elements T_r. As the set of targets is unknown (it is the task of the correlation algorithm to make the best possible estimation of the set), T may contain both firmly established objects from before, and new potential targets to be considered. An example of track and target numbering is shown below in Fig. 1.

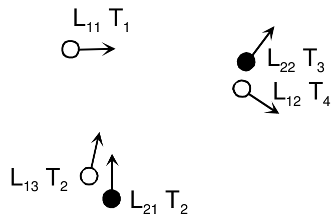

Fig. 1. An example on how target numbers relate to track numbers.

Source 1 provides local tracks L_{11}, L_{12} and L_{13}, and source 2 has L_{21} and L_{22}. The algorithm has assigned target T_1 to L_{11}, T_2 to L_{13} and L_{21}, T_3 to L_{22}, and T_4 to L_{12}. Here thus L_{13} and L_{21} are considered to represent the same target.

The above is just one example of correlation; for example, the algorithm might have concluded that L_{13} and L_{21} should represent two different targets or that L_{12} and L_{22} should be the same, or that T_3 and T_4 should be swapped. Each set of assignments represents one correlation hypothesis, and each hypothesis assigns the targets to zero or more of the local tracks (zero meaning that the target does not exist or is not tracked). Any assignment is permitted, except that no target may be assigned to more than one local track from the same source.

2.2 Hypotheses

The hypotheses can conveniently be represented by a table as follows (this example is unrelated to Fig. 1):

Table 1. A representation of correlation hypotheses

T_1	T_2	T_3
L_{12}	L_{11}	
	L_{21}	
	L_{32}	L_{31}

Here source 2 provides one track, while sources 1 and 3 provide two tracks each. Further it is assumed that there are two targets from before, and that we allow at most one new target. In the hypothesis shown in the table, the first track from the third source, L_{31}, does not go with any of the of the hitherto handled targets, but instead represents either a false track or a new real target. One interpretation is that there were in fact three targets, but due to poor sensor resolution, no single sensor has yet detected all three. The number of possible hypotheses with local tracks as above, and two old plus at most one new target, is $6 \cdot 3 \cdot 6 = 108$, and it is up to the correlation algorithm to decide which one is the most likely. The number of hypotheses explodes with the number of tracks and targets.

In forming hypotheses, note that if we allow more than one new target, then some apparently different hypotheses will be the same. If there were a T_4 column in the table above, then putting L_{31} under T_4 instead of T_3 would not be a new hypothesis.

The key idea behind the correlation and numbering algorithm is that we consider a large number of correlation hypotheses $H_n, n = 1 \ldots$, where each hypothesis stands for a possible assignment of targets (or target numbers) T_r to local tracks L_{si}, and that we estimate the probability $p(H_n)$ of each hypothesis recursively at prediction as well as updating steps.

2.3 Prediction step

The prediction step amounts to computing the probabilities $p(H_n, t_k^-)$ from earlier values $p(H_n, t_{k-1}^+)$, taking into account that mistakes like track swapping or track loss could occur in the meantime. The probability mass tends to go from the most likely hypothesis to other, less likely, hypotheses, as there is always some risk that the previous assignments do not hold because of tracking errors. Especially under difficult conditions, such as dense target situations, inadequate sensor coverage, sparse measurements, etc., there is an increased probability for tracking errors. Expressing these probabilities is however very difficult. Some simplification can be made along the lines described in Sec. 4.

2.4 Updating step

The updating step is preferably (but not necessarily) based on Bayes' rule, according to which

$$p\left(H_n, t_k^+\right) \propto p\left(z \middle| H_n\right) p\left(H_n, t_k^-\right) \tag{1}$$

where $p\left(z \middle| H_n\right)$ is the probability of obtaining the observed tracks and their observed positions etc. assuming that H_n is true. However, in order for Bayes' rule to be applicable in the form given above it is important that the sequence of observations is independent. Handling dependent observations in a Bayesian framework is complicated and requires accurate knowledge of the nature of dependencies. We avoid this problem by not updating too frequently, e.g., the tracks must have moved a long way compared to the positional uncertainties involved. Failing to do so can have the result that one of the hypotheses will incorrectly dominate over the others.

For simplicity we will here base the calculation on track positions only; extensions to include more data like velocities and other attributes are relatively straightforward.

Each hypothesis H_n can be represented as in Table 1. It is defined by a set T of targets, each corresponding to a target number T_r and (unknown) true positions x_r. To each target corresponds a set L_{nr} of local tracks L_{sj_r}, with indices indicating that the track is the j^{th} from sensor s corresponding to target r. Each track has the position

z_{sj_r}. (For example, in Table 1, the set L_{n2} consists of L_{11}, L_{21} and L_{32}). The positions are described by probability densities $f_{sj}(z,x)$ around the true positions. The probability of finding the track within a small volume dz around z_{sj_r} is $f_{sj_r}(z_{sj_r}, x_r)\,dz$. The probability of finding all the tracks corresponding to all targets at their given positions is then

$$\prod_{T_r}\prod_{L_{sj_r}} f_{sj_r}\left(z_{sj_r}, x_r\right)(dz)^N \qquad (2)$$

where N is the total number of local tracks in the cluster. The true positions x_r are not known at all. They could be found in any volume element dx_r with corresponding probability $\rho\,dx_r$ where ρ us a typical target density (probably a tuning parameter). We have to integrate over all possible values, that is,

$$p\left(z\middle|H_n\right) \propto \rho^{M_n}\prod_{T_r}\int\prod_{L_{sj_r}} f_{sj_r}\left(z_{sj_r}, x_r\right)dx_r \qquad (3)$$

where M_n is the number of targets in this hypothesis, and where the factors containing dz are left out as they are the same for all hypotheses.

The usage of probability densities in this way is an approximation as one neglects the dependence between the probability densities stemming from the same target (e.g. if the target makes a sudden maneuver, the error tends to have the same sign in all tracks representing this target). There are techniques in the literature to improve the correlation accuracy in this respect, see e.g. [1], Sec. 9.6.3, and references therein. Such methods can also be included in the Bayesian track correlation described here.

There is however a complicating factor that should not be neglected, and that is visibility. Each sensor may have a limited capability to see a certain target. The sensors may have different coverage areas in space or electromagnetic spectrum, and thus it is perfectly possible that a target is not seen by some of the sensors, although it is within the sensors' nominal range. As a result, the algorithm could for example force a ground target track to be correlated with an air target track. The probability that a target T_r is seen by a source s will be expressed by a visibility factor $0 \le K_{sr} \le 1$. Likewise, the probability that the target is not seen by the source is $1-K_{sr}$. As an example, for the hypothesis represented by Table 1, we should include visibility factors according to

Table 2. Visibility factors

T_1	T_2	T_3
K_{11}	K_{12}	$1\text{-}K_{13}$
$1\text{-}K_{21}$	K_{22}	$1\text{-}K_{23}$
$1\text{-}K_{31}$	K_{32}	K_{33}

Evaluating the probabilities for the hypotheses according to visibility then yields

$$p\left(z\middle|H_n\right) \propto \prod_{r=1}^{M_n}\prod_{s=1}^{N_s} m\left(K_{sr}\right) \qquad (4)$$

where

$$m\left(K_{sr}\right) = \begin{cases} K_{sr}, \text{source } s \text{ has target } r \\ 1-K_{sr}, \text{source } s \text{ does not have target } r \end{cases} \qquad (5)$$

and N_s is the number of sources involved.

One should however be cautious when using the visibility factors in (4) as (1) is meant to be used recursively, and the visibility factors will not change rapidly. As a result the same information would be entered over and over again into the recursion loop, and this is of course not correct, unless the visibility factors are 0 or 1. (Actually using only 0 or 1 for the visibility is an easy albeit suboptimal way out of this problem. Another way could be keeping the visibility out of the recursion loop, and only applying it at the output.) Similar arguments hold if target classification or other attributes are incorporated in the application of Bayes rule. Using classification information is important to prevent correlating e.g. friend and foe.

In practice one would usually assume Gaussian distributions for the probability densities. Then the uncertainty of a track is characterized by its covariance matrix P_{sj_r}, and the probability density is given by

$$f_{sj_r}\left(z_{sj_r}, x_r\right) = \frac{1}{\sqrt{(2\pi)^D \det P_{sj_r}}}\, e^{-\frac{1}{2}\left(z_{sj_r}-x_r\right)^T P_{sr}^{-1}\left(z_{sj_r}-x_r\right)} \qquad (6)$$

where D is the number of dimensions. Then the integration in (3) can be carried out explicitly, leading to a somewhat complicated, but definitely useable, expression.

3. Simplified representation of hypotheses

The number of hypotheses that can represent a cluster may be very large and it can be very difficult to manage all the corresponding hypotheses. We will therefore introduce a significantly simplified representation, where *each local track and each target is handled individually*. The steps are explained in the following sections, and show how the probabilities for the hypotheses can be computed from the simplified representation.

3.1 Characterization of tracks

In the simplified representation each track L_{si} is characterized by its most likely target T_{si} and the probability p_{si} that this is correct (the f_{sj_r} of Sec. 2.4 does give the probability density assuming that we know the corresponding target, but we can't generally be certain that the assumption is correct). The probability that any arbitrary target T_a in the cluster is the correct one for L_{si} is then given by

$$p(L_{si}, T_a) = \begin{cases} p_{si}, T_a = T_{si} \\ \dfrac{1 - p_{si}}{M_n}, T_a \neq T_{si} \end{cases} \quad (7)$$

Of course one can improve on this as all alternative targets need not be equally likely.

Each correlation hypothesis H_n stands for an assignment $T_n(L_{si})$ of the targets to tracks. The probability for the hypothesis based on the assignments can thus be written

$$p^L(H_n) = \prod_{L_{si}} p(L_{si}, T_n(L_{si})) \quad (8)$$

3.2 Characterization of targets

Each target T_r is characterized by the probability p_r^T that this target indeed exists and is trackable (maybe it has run out of sensor coverage, or ceased to exist, or never existed). Introduce the function

$$p^T(T_r) = \begin{cases} p_r^T, T_r \text{ is used in hypothesis} \\ 1 - p_r^T, T_r \text{ is not used in hypothesis} \end{cases} \quad (9)$$

Each hypothesis is characterized by the set of targets involved. The selection of targets will then have the probability

$$p^T(H_n) = \prod_r p^T(T_r) \quad (10)$$

3.3 Probabilities for the hypotheses

To compute the probability of a hypothesis H_n from the simplified representation one has to

i. Note the selection of targets in the hypothesis
ii. Compute the probability factors $p^L(H_n)$ for all hypotheses consistent with the selection of targets
iii. Compute the approximate conditional probability

$$p^C(H_n) = \frac{p^L(H_n)}{\sum_m p^L(H_m)} \quad (11)$$

where the sum runs over all hypotheses consistent with the selection of targets of H_n

iv. Compute the probability for H_n from

$$p(H_n) = p^C(H_n)\, p^T(H_n) \quad (12)$$

The whole procedure is illustrated by a numerical example in the Appendix.

3.4 From full to simplified representation

We have now illustrated how to go from the simplified representation (i.e. data are stored individually for each track and for each target) to the full (i.e. the track/target tables.) We must also be able to go back to the simplified representation. This is in principle very straightforward. First make sure that the hypotheses are properly normalized. Pick the target numbers that are assigned according to the most likely hypothesis, and then add the probabilities of all hypothesis assigning the target number in question to the track. Likewise, add the probabilities of all hypotheses where the first target number is included, etc. Please refer to the Appendix for a numerical example.

A more sophisticated version of the simplified representation could consist in storing more than one target number and its probability in each track. Although this could be very helpful in allowing secondary hypotheses to grow, it adds significantly to the complexity of the algorithm.

4. Summary

The complete processing loop, including the optional simplified representation, is shown in Fig. 2.

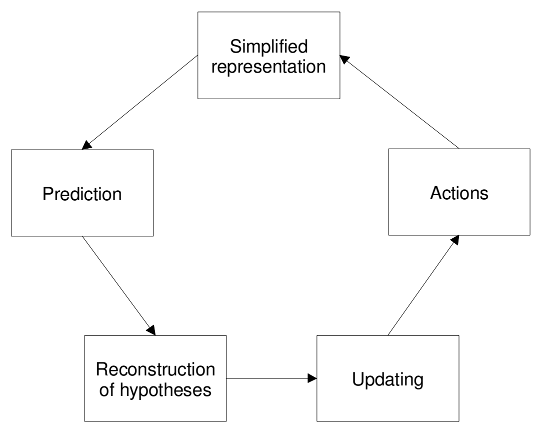

Fig. 2. The main steps in the processing loop.

Each local track carries information about the most likely target and the probability that the assignment is correct, and there is also a list maintaining the targets and the probability that the target in fact exists.

The prediction step is very difficult to express other than empirically. Ideally each tracker should give an estimate for each track that the track/target linkage is broken. This could e.g. be expressed as a probability per time unit, and should depend on e.g. on target density, false alarm density, sensor coverage, sensor characteristics.

At the reconstruction step, the hypotheses are (in principle) expressed as track/target tables, such as Table 1, and the probabilities for the different hypotheses are computed from (12). There must also be allowance for new targets to appear.

At updating, new probabilities for the hypotheses are computed using Bayes' rule, that is, Eqs. (1, 3). Occasionally also visibility (Eq. (4)) and other attributes can be included.

If the number of hypotheses is large one should not try to evaluate every one. It is sufficient to look at a subset of hypotheses, namely those that a naïve correlation algorithm would produce, and those hypotheses that can be formed from the naïve by a limited number of steps such as swappings or changing a single track/target assignment. One may also similarly start with the hypothesis from the previous cycle and evaluate a limited number of deviations from that.

The resulting most likely correlation hypothesis may deviate from the currently valid, and then actions should be taken, with or without operator approval. These actions consist of effecting track swapping, track split/merge, removal of target, creation of target, etc.

Finally, the simplified representation is recomputed as described in Sec. 3.4 and stored in the tracks and target list.

5. Conclusion

We have presented a technique for combined track-to-track correlation and track numbering. Although one cannot say there is any standard way of doing the task, the approach here has many advantages over what is usual:

- Sound Bayesian evaluation of different correlation hypotheses
- Continual updating and recorrelation when needed
- Allows many factors to be taken into account which really matter, such as visibility and probability of tracking mistakes

Admittedly some of the input parameters may be difficult to obtain, and admittedly the approximations when using the simplified representation, are quite rough – but we still consider this method as superior to those that cannot consider the complicating factors at all.

Unfortunately, the impacts of the approximations have not yet been fully analyzed. Nor have we made any comparison of the performance with standard methods, much because we are not aware of any such standard method for track numbering. However, a prototype has been implemented and tested with live data in connection with a system for distributed correlation, i.e. correlation simultaneously taking place in different nodes over a national network. The results so far look very good.

Appendix: A numerical example

To clarify the transitions between the simplified representation and the full, we here give a numerical example. The notation is taken from Sec. 3.

Suppose we have one track L_{11} from source 1, and one track L_{21} from source 2. They have been assigned different target numbers, T_1 and T_2, so that

$$T_{11} = T_1, p_{11} = 0.8$$
$$T_{21} = T_2, p_{21} = 0.7$$

Moreover, we are rather sure that T_1 is a genuine target, while we are less certain about T_2:

$$p_1^T = 0.99$$
$$p_2^T = 0.6$$

We want to know if the tracks should have the same target number, and if so, which number, or if they should still be considered to represent different targets. The hypotheses are H_1 (no change):

T_1	T_2
L_{11}	
	L_{21}

H_2, both are T_1:

T_1	T_2
L_{11}	
L_{21}	

H_3, both are T_2:

T_1	T_2
	L_{11}
	L_{21}

and the unrealistic H_4, where both target numbers are replaced:

T_1	T_2
	L_{11}
L_{21}	

From the previously given equations we compute

$$p^T(H_1) = p^T(H_4) = 0.99 \cdot 0.6 = 0.594$$
$$p^T(H_2) = 0.99 \cdot 0.4 = 0.396$$
$$p^T(H_3) = 0.01 \cdot 0.6 = 0.006$$

Note 1: With respect to target selection, H_1 and H_4 are the same hypothesis, therefore the probabilities as written here add up to more than 1.

Note 2: As the selection "Neither T_1 nor T_2" is not possible here (as we have chosen not to include a T_3 column) the probabilities for the three selections add up to slightly less than one. Strictly the probabilities should be renormalized so that their sum is one, but since what we need is only to find the most probable hypothesis this is not needed.

Next, compute the probability factors $p^L(H_n)$

$$p^L(H_1) = 0.8 \cdot 0.7 = 0.5$$
$$p^L(H_2) = 0.8 \cdot 0.3 = 0.24$$
$$p^L(H_3) = 0.2 \cdot 0.7 = 0.14$$
$$p^L(H_4) = 0.2 \cdot 0.3 = 0.06$$

and the conditional probabilities

$$p^C(H_1) = \frac{0.56}{0.56 + 0.06} = 0.903$$
$$p^C(H_2) = \frac{0.24}{0.24} = 1$$
$$p^C(H_3) = \frac{0.14}{0.14} = 1$$
$$p^C(H_4) = \frac{0.06}{0.56 + 0.06} = 0.097$$

and finally the resulting probabilities

$$p(H_1) = 0.903 \cdot 0.594 = 0.536$$
$$p(H_1) = 0.396$$
$$p(H_3) = 0.006$$
$$p(H_4) = 0.097 \cdot 0.594 = 0.058$$

Thus the first hypothesis is the most likely and the tracks should not be merged.

Next we will show how to go back to the simplified representation. In the example given above, the normalized probabilities are

$$p(H_1) = 0.538$$
$$p(H_1) = 0.398$$
$$p(H_3) = 0.006$$
$$p(H_4) = 0.058$$

The first hypothesis is the most likely one, thus

$$T_{11} = T_1$$
$$T_{21} = T_2$$

as before, and

$$p_{11} = 0.511 + 0.377 = 0.936$$
$$p_{21} = 0.511 + 0.057 = 0.544$$
$$p_1^T = 0.511 + 0.377 + 0.055 = 0.994$$
$$p_2^T = 0.511 + 0.057 + 0.055 = 0.602$$

We did not get back the original values, but that should not have been expected. In general a very large number of probabilities in the complete set are represented as a smaller number of probabilities; this is a process that cannot be reversed.

References

[1] Samuel Blackman and Robert Popoli. Multiple sensor tracking: Issues and Methods in *Design and Analysis of Modern Tracking Systems*, Pages 595 - 637, Artech House, Boston, London, 1999.

[2] Oliver E. Drummond. A Hybrid Sensor Fusion Algorithm Architecture and Tracklets in *Proceedings of SPIE, Signal and Data Processing of Small Targets 1997*, Vol 3163, pages 485-502, San Diego, CA, USA, 29-31 July 1997.

[3] D. J. Percival and K.A.B. White. Multihypothesis Fusion of Multipath Over-the-horizon Radar Tracks in *Proceedings of SPIE, Signal and Data Processing of Small Targets 1998*, Vol 3373, pages 440-451, Orlando, FL, USA, 14-16 April 1998.