A Transferable Belief Model Approach to Combat Identification

William J. Farrell III

ABSTRACT

Combat identification (CID) is the process of accurately characterizing battlespace entities to enable high-confidence, real-time application of tactical options, such as engagement. Evidence to support CID estimates is often sparse, latent in the battlefield, or both, raising the risk of association ambiguity and potential loss of CID custody. Therefore, an automated CID estimation methodology must properly account for and convey its results' uncertainty, ambiguity, and ignorance to the warfighter to support timely, well-informed decision-making. The automated CID estimation process presented in this article is a computationally scalable approach to achieve robust CID custody in over-the-horizon targeting applications. Novel aspects of this approach include (1) a compact representation of track histories as tracking segments (vice measurements); (2) a temporal history of kinematic ambiguities between tracks; and (3) a transferable belief model for open-world evidential reasoning under uncertainty, ambiguity, and conflict. The result is an actionable, informative CID estimation process that accounts for real-world challenges and constraints.

INTRODUCTION AND BACKGROUND

Combat Identification Evidence for Targeting Decisions

The targeting process requires the warfighter to identify a battlespace entity, at a confidence level determined by rules of engagement, in order to make an engagement decision. This combat identification (CID) process often requires the warfighter to manually assess evidence accrued on the target. Because of the sparsity and/or latency of highly informative CID evidence (e.g., imagery), it may be difficult to uniquely associate this evidence with tracked battlespace objects. Additionally, when tracks of battlespace objects experience gaps or become closely spaced, it is unclear whether previously associated CID evidence can be attributed to a potential target at the time of an engagement decision. A credible and trustworthy automated CID estimation decision aid must address these challenges.

This article presents a mathematically rigorous approach for producing actionable CID estimates in the presence of uncertainty, ambiguity, conflict, and an open-world assumption. Specifically, this article presents an evidence fusion and inference methodology that provides actionable track identification estimates using associated CID evidence from multiple heterogeneous

intelligence (INT) sources. This approach accounts for the four key challenges of information fusion:

1. Imprecise CID evidence about battlespace objects (uncertainty)
2. Imprecise association of CID evidence with a single battlespace object (ambiguity)
3. Associated CID evidence from multiple sources that may not be completely consistent (conflict)
4. Ignorance of the universe of possible target identity hypotheses (open-world assumption)

In addition to addressing these four challenges, the methodology offers features that make for a practical implementation of CID estimation:

- The ability to reason over a taxonomy of possible target identity hypotheses in lieu of a simple set of hypotheses (hierarchical reasoning)
- The ability to adapt the reported CID estimate to an appropriate level of a hypothesis taxonomy (actionable)
- The ability to incorporate evidence from multiple heterogeneous INT sources whose feature spaces do not overlap (multi-INT)

The CID estimation approach is then extended to accommodate latent CID evidence by (1) introducing a computationally scalable method for latent data association of evidence with tracks and (2) combining CID evidence to account for tracking ambiguities (e.g., closely spaced tracks) that may have occurred after the time of evidence association. The latter capability is critically important to the warfighter. Imagine a scenario where CID evidence is associated with an adversary track that subsequently becomes ambiguous with the track of a civilian (i.e., neutral) platform. If the warfighter is not aware that an ambiguity occurred after associating evidence, their assessment of the CID evidence will not have proper context, which could lead to a catastrophically erroneous decision. Thus, any CID estimation decision aid must properly account for tracking ambiguities occurring after CID evidence is associated.

CID Evidence for Improved Track Custody

In addition to fulfilling the primary need to accurately estimate CID when information is latent, CID evidence can also be used to enhance track custody in the presence of tracking gaps (track stitching) and to reestablish CID after periods of tracking ambiguity (CID disambiguation).

Track stitching is the process of associating newly formed tracks with previously dropped tracks. It is difficult to stitch tracks using only kinematic comparisons when there are significant tracking gaps, but comparing CID estimates to eliminate incompatible tracks can improve the process. The benefit of track stitching is that once the new track is stitched, it can inherit important previously derived contextual information from the stitched dropped track, such as intent, point of origin, pattern of life assessments, and other historical attributes that may support an engagement decision. The process of track stitching is conceptually straightforward. First, newly established tracks are evaluated against dropped tracks for kinematic feasibility to obtain candidates for stitching. Second, once CID evidence is associated with the new track, this evidence is compared to the CID estimate for all candidate stitched tracks obtained from the kinematic feasibility assessment. With sufficiently consistent CID evidence, tracks may be determined to represent the same battlespace object. A statistical hypothesis test can help the warfighter make a decision to stitch tracks.

When tracks become ambiguous, a trustworthy CID estimation algorithm must acknowledge that it can no longer determine which track holds which CID estimate, and the fused, ambiguous CID estimate must be attributed to both ambiguous tracks (CID entanglement). However, CID evidence that is unambiguously associated with a track after an ambiguity has occurred can be used to estimate the CID of the associated track as well as the CID of tracks that were entangled with it. Consider an example of two crossing tracks, each with unique unambiguous CID estimates (adversary and civilian, respectively) before they cross. As the tracks cross, their CID estimates become entangled, and it is not clear which track is an adversary and which is a civilian after they separate. Once one of the track's associated CID evidence indicates that it is a civilian platform, we can deduce that the other track must be an adversary. An automated system can make this powerful deduction on behalf of the warfighter to shorten the targeting decision timeline.

Article Organization

This article is broken into two main sections: (1) CID estimation and (2) latent CID evidence association.

In the CID Estimation section, first the core CID estimation mathematical framework, the transferable belief model (TBM), is justified and presented. Second, the TBM mathematics are applied to a taxonomy of classification hypotheses, and example calculations are presented. Third, a computationally efficient TBM implementation is presented using a BitSet encoding of the hypothesis taxonomy. Fourth, a methodology for defining multi-INT CID features and evidence evaluation is presented. Last, an information theoretic algorithm is derived to determine the best level to report the CID estimate within a taxonomy of hypotheses.

The second major section addresses the specific challenges of associating latent CID evidence with historical tracks in a computationally scalable way. First it presents the concept of the track segment graph (TSG)

for an efficient and compact representation of historical track kinematics and periods of ambiguity. Then it presents the use of an interval binary search tree (IBST) for retrospective search of candidate track associations. The section concludes with details on an algorithm for propagating CID estimates from the time of associated evidence to the current time while accounting for ambiguities that may have occurred.

CID ESTIMATION

Problem Statement

This article considers CID estimation as a special case of the more general problem of estimating the belief distribution across a large set of discrete hypotheses representing the type of an object being observed in the battlespace. More generally, this hypothesis space can be hierarchical and represented as a taxonomy of possible object types. Real-world examples of taxonomies related to targeting include those defined by Link 16,[1] MIL-STD-2525D,[2] and Cursor on Target (CoT).[3] The Link 16 taxonomy represents battlespace objects at three hierarchical levels: category (air, land, surface, etc.), platform (carrier, destroyer, fighter, etc.), and specific type (CVN, DDG, F-18, F-16, etc.). The MIL-STD-2525D taxonomy was defined for display symbology, and its hierarchical depth varies across the taxonomy. Finally, the CoT taxonomy is similar to that of MIL-STD-2525D and is used in the examples throughout this article.

In addition to supporting a hierarchical hypothesis space, the CID estimation approach must provide a framework for reasoning under uncertainty, ambiguity, conflict, and an open-world assumption. Uncertainty is due to imprecise CID evidence that reflects ignorance about a battlespace object's observed features. Ambiguity is the result of imprecise association or attribution of evidence to a single battlespace object. Conflict can arise when different sensing modalities (multiple INTs) have different degrees of ignorance or there is association ambiguity, which is common when reasoning over sensing modalities with heterogeneous feature spaces. Finally, the open-world assumption accommodates the reality that the hypothesis space may be incomplete.

The TBM

The two most commonly used mathematical frameworks for reasoning under uncertainty are Bayesian inference[4] and Dempster–Shafer (D-S) evidence combination.[5] However, neither framework employs the open-world assumption, as indicated by their renormalization of probability/belief over the defined hypothesis space. As a result, these frameworks assign probability/belief to the hypotheses that are most consistent with the evidence, even if no hypothesis is significantly consistent, which can be misleading to a decision-maker.

To illustrate the value of employing the open-world assumption, consider the so-called Zadeh counterexample problem.[6] A patient is experiencing headaches and visits a doctor who assesses that the symptoms are 99% likely the result of a migraine and 1% likely the result of a stroke. Seeking a second opinion, the patient visits a second doctor who assesses that the symptoms are 99% likely the result of allergies and 1% likely the result of a stroke. Using Bayesian inference or D-S evidence combination, combining the evidence from both doctors results in a definitive explanation (with 100% likelihood) that the patient is experiencing a stroke (a conclusion that both doctors found highly unlikely, at only 1%). Alternatively, an evidence combination approach called the transferable belief model (TBM), first published by Smets in 1990,[7] declares that it is 0.01% likely that the patient is experiencing a stroke and 99.99% likely that the patient is suffering from something else that neither doctor considered as a possibility. Both doctors are reasoning over the same hypothesis space (i.e., migraine, allergies, and stroke); however, this hypothesis space is incomplete. Therefore, Zadeh's counterexample exposes the inability of Bayesian and D-S frameworks to deal with ignorance of the complete hypothesis space. The TBM's explanation of the patient's encounter with both doctors is more aligned with a common-sense inference and justifies employing the TBM evidence combination approach for CID estimation.

To formalize the TBM framework, the following definitions and notations are introduced:

- **Frame of discernment (FOD):** Denoted by Ω, this is a finite set of mutually exclusive elements of a hypothesis space. Because of the open-world assumption, the FOD may not be exhaustive.
- **Basic belief assignment (BBA):** A mapping m^{Ω} from the power set 2^{Ω} over the FOD to a value between 0 and 1, $2^{\Omega} \longrightarrow [0,1]$, that satisfies

$$\sum_{A \in 2^{\Omega}} m^{\Omega}(A) = 1. \quad (1)$$

- **Focal elements:** Elements $A \in 2^{\Omega}$ where $m^{\Omega}(A) > 0$.
- **Vacuous belief function (VBF):** The belief function that represents total ignorance over the hypothesis space, $m^{\Omega}(\Omega) = 1$.

Of particular note, the VBF is reminiscent of the so-called diffuse prior used in Bayesian inference. However, there is a key conceptual difference. In the Bayesian framework, a diffuse prior yields an initial probability of $1/N$ across a hypothesis space of size N. Within the TBM framework, a vacuous prior is a set theoretic statement assigning all the initial belief to the entire set of hypotheses with no assertion of how it is divided among that set. In contrast to D-S, the TBM allows for the empty set

to be a focal element (i.e., to have non-zero belief mass explicitly representing the open-world assumption).

The TBM framework defines evidence combination rules for conjunctive (logical AND) and disjunctive (logical OR) reasoning. Within the TBM framework,[6] the conjunctive rule is based on the assumption that the belief functions to be combined are induced by reliable sources of information, whereas the disjunctive rule only assumes that at least one source of information is reliable, but we do not know which one. Within the context of tracking systems, conjunction is used when we have high confidence that the CID evidence to be combined has been associated with the correct track, whereas disjunction is used when there is ambiguity as to which CID evidence is associated with a track. Mathematically, these combination rules are as follows[8]:

- **Conjunctive combination rule** for two BBAs, 1 and 2, over the same FOD:

$$m^{\Omega}_{1\cap 2}(A) = \sum_{X\cap Y=A} m^{\Omega}_1(X)\, m^{\Omega}_2(Y) \quad \forall A \neq \emptyset \in 2^{\Omega}\ (X, Y \in 2^{\Omega}) \tag{2}$$

$$m^{\Omega}_{1\cap 2}(\emptyset) = 1 - \sum_{A\in 2^{\Omega}, A\neq\emptyset} m^{\Omega}_{1\cap 2}(A)$$

- **Disjunctive combination rule** for two BBAs, 1 and 2, over the same FOD Ω:

$$m^{\Omega}_{1\cup 2}(A) = \sum_{X\cup Y=A} m^{\Omega}_1(X)\, m^{\Omega}_2(Y) \quad \forall A \neq \emptyset \in 2^{\Omega}\ (X, Y \in 2^{\Omega}) \tag{3}$$

$$m^{\Omega}_{1\cup 2}(\emptyset) = 1 - \sum_{A\in 2^{\Omega}, A\neq\emptyset} m^{\Omega}_{1\cup 2}(A)$$

The conjunctive combination rule creates a BBA over the subsets of 2^{Ω} that result from the intersection of the BBAs 1 and 2. The value of belief assigned to each of these subsets A is given by the sum of the products of the BBAs 1 and 2 where an intersection results in A. A similar interpretation using a set union is applicable to the disjunctive combination rule. One key distinction between the TBM and D-S is that the TBM does not renormalize over conflict between combined belief masses and explicitly assigns belief mass to the empty set equal to the conflict between the combined belief masses.

The D-S conjunctive combination rule does not assign mass to the empty set and instead renormalizes the beliefs over the FOD as follows[5]:

$$m^{\Omega}_{1\oplus 2}(A) = \begin{cases} \frac{\sum_{X\cap Y=A} m^{\Omega}_1(X)\, m^{\Omega}_2(Y)}{1-\sum_{X\cap Y=\emptyset} m^{\Omega}_1(X)\, m^{\Omega}_2(Y)} & \forall A \in 2^{\Omega},\ A \neq \emptyset, (X, Y \in 2^{\Omega}) \\ 0 & A = \emptyset \end{cases}. \tag{4}$$

The final piece of the TBM framework, which differs from the D-S approach, is the decision-making procedure. In the TBM, decisions are not made based on BBAs resulting from Eqs. 2 and 3, but rather from a probability transformation called the pignistic probability transformation (PPT), which Smets defines as[6]

$$BetP^{\Omega}(A) = \sum_{B\in 2^{\Omega}} \frac{|A\cap B|}{|B|} \frac{m^{\Omega}(B)}{1-m^{\Omega}(\emptyset)} \qquad \forall A \subseteq \Omega\,, \tag{5}$$

where $|X|$ is the number of elements of Ω in set $X \in 2^{\Omega}$.

The Zadeh Counterexample

Returning to the Zadeh counterexample, a comparison between D-S and the TBM can be presented in detail by evaluating Eqs. 5 and 2 for the Zadeh problem presented above. Both D-S and the TBM use BBAs to express evidence, which in the Zadeh problem is given by, excluding all zero masses, the following[6]:

$$m^{\Omega}_1(a) = 0.99,\ m^{\Omega}_1(b) = 0.01$$
$$m^{\Omega}_2(c) = 0.99,\ m^{\Omega}_2(b) = 0.01\text{'}$$

where 1 and 2 subscripts are used for the evidence obtained from the first and second doctor, and we have the FOD (diagnoses) a = migraine, b = stroke, c = allergies. For D-S, the result of combining these two BBAs is

$$m^{\Omega}_{1\oplus 2}(b) = \frac{\sum_{X\cap Y=b} m^{\Omega}_1(X)\, m^{\Omega}_2(Y)}{1-\sum_{X\cap Y=\emptyset} m^{\Omega}_1(X)\, m^{\Omega}_2(Y)} = \frac{(.01)(.01)}{1-((.99)(.01)+(.99)(.99)+(.99)(.01))} = \frac{.0001}{1-.9999} = \frac{.0001}{.0001} = 1\,.$$

This indicates complete confidence in the patient experiencing a stroke. Alternatively, the TBM result for combining these two BBAs is

$$m^{\Omega}_{1\cap 2}(b) = \sum_{X\cap Y=b} m^{\Omega}_1(X)\, m^{\Omega}_2(Y) = (.01)(.01) = .0001$$

$$m^{\Omega}_{1\cap 2}(\emptyset) = 1 - \sum_{A\in\neq\emptyset 2^{\Omega}} m^{\Omega}(A) = .9999,$$

which yields the more sensible result that there is a high confidence that the true diagnosis is not within the set of migraine, stroke, or allergies. That is, mathematically, the true diagnosis is in the empty set ø with 99.99% belief.

The TBM with a Taxonomy of Hypotheses

Closer inspection of the TBM Eqs. 2 and 3 (and indeed the D-S combination rules) reveals that evidence combination, apart from simple multiplication and summation, reduces to set operations of intersections and unions. Comparatively, a taxonomy is, by definition, structured in a way that makes these operations

trivial. For example, consider the classification taxonomy in Figure 1 with a three-level hierarchy (e.g., the Link 16 standard taxonomy), a root node representing the set Ω, and the open-world assumption modeled by an empty set.

This example illustrates the following relevant facts regarding taxonomies:

- Each node is precisely given by the union of all its children nodes. For example, $t_2^1 = \{t_3^2, t_4^2, t_5^2\}$.
- The specificity of the hypotheses increases as the taxonomy is traversed from top to bottom. For example, hypotheses at the level T^2 are more specific hypotheses than their parent nodes at the level T^1.
- Nodes are mutually exclusive (there is no overlap in their hypotheses) if they are on different branches of the taxonomy.

And thus, the FOD is the set of all leaf nodes (nodes with no children).

The TBM can be used to combine evidence from different sources in the form of BBAs that express evidence on any subsets of the hypotheses in the taxonomy, including at different levels in the hierarchy. For example, one source of evidence (e.g., sensor) may be particularly good at distinguishing between classes of battlespace objects enumerated in level T^1, while another source may exploit phenomenology that is particularly good at identifying classes of battlespace objects that are contained in the sub-tree of t_2^1, and still a third source may only be able to detect and discriminate between classes of battlespace objects t_1^2 and t_2^2. This ability to combine evidence that spans different subsets of the hypothesis space enables the use of heterogeneous evidence sources reporting on heterogeneous sets of hypotheses. This is fundamentally

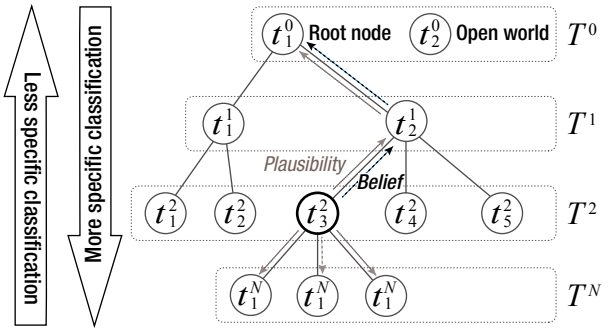

Figure 1. Simple example hypothesis taxonomy. This example taxonomy has three levels, not counting the root level, where the third level is denoted T^N. Each element in the taxonomy is denoted t_j^i where i indicates the level, and j indicates the node's position within that level. This also illustrates how the mass assigned to node t_3^2 propagates as belief (dark arrows) and plausibility (light arrows) through the taxonomy to each level of specificity in the taxonomy. (Copyright 2012 Society of Photo-Optical Instrumentation Engineers [SPIE].[9])

what enables the TBM estimation approach to support multi-INT fusion of evidence for CID.

Once TBM combination rules have been applied, the algorithm has to make a decision about what to report. As discussed by Smets,[6] the resulting fused BBAs are not suitable for making a decision because their masses are, in general, scattered throughout the taxonomy, across hierarchical levels. An alternative often used in classifier systems is to simply report every hypothesis that has some significant belief along with their belief value. This is not a desirable product for a potentially burdened warfighter who may have to evaluate CID estimates quickly to support a decision. An alternative, and more actionable, approach is to report the most concise CID while limiting the loss of specificity. For example, instead of reporting a CID estimate of t_3^2, t_4^2, or t_5^2, it is more useful and readily comprehendible to simply report t_2^1, so long as the loss of specificity is not too great. The appropriate reporting level can be achieved using the formal definition of a taxonomy and applying the *BetP* PPT. Specifically, a suggested methodology is to start at the lowest level (i.e., the set of leaf nodes) in the taxonomy and compute the *BetP* using the fused BBA from all CID evidence. The *BetP* is then computed at one level higher in the taxonomy, and the information loss between the two levels is computed. This process repeats until choosing a level in the taxonomy would exceed the specified level of specificity loss. The Adaptive Output section details an information theoretic method for selecting the taxonomy level at which to report the CID estimate.

CID Estimation Rules

Now, using the hypothesis taxonomy in Figure 2, consider a situation resulting in CID entanglement between two crossing tracks. The first track, $\mathcal{T}1$, has a BBA, $m_{\mathcal{T}1(\tau = 0)}$, over this hypothesis taxonomy (omitting the Ω superscript for brevity) given by the CID evidence unambiguously associated with it before the crossing at time $\tau = 0$. Likewise, the second track, $\mathcal{T}2$, has a BBA, $m_{\mathcal{T}2(\tau = 0)}$, over the same hypothesis taxonomy. As an example, the BBAs are

$$\begin{matrix} m_{\mathcal{T}1(\tau=0)}\{t_1^2\} = .9, m_{\mathcal{T}1(\tau=0)}\{t_2^2\} = .1, \\ m_{\mathcal{T}2(\tau=0)}\{t_1^2\} = .1, m_{\mathcal{T}2(\tau=0)}\{t_2^2\} = .9 \end{matrix} \quad (6)$$

which trivially yields a *BetP* for $\mathcal{T}1$ of

$$BetP_{\mathcal{T}1(\tau=0)}(\{t_1^2\}) = 0.9, BetP_{\mathcal{T}1(\tau=0)}(\{t_2^2\}) = 0.1 \quad (7)$$

and a *BetP* for $\mathcal{T}2$ of

$$BetP_{\mathcal{T}2(\tau=0)}(\{t_1^2\}) = 0.1, BetP_{\mathcal{T}2(\tau=0)}(\{t_2^2\}) = 0.9. \quad (8)$$

Thus, it is easy to see that the CID estimates of these tracks might be reported at the taxonomy level T^2 since

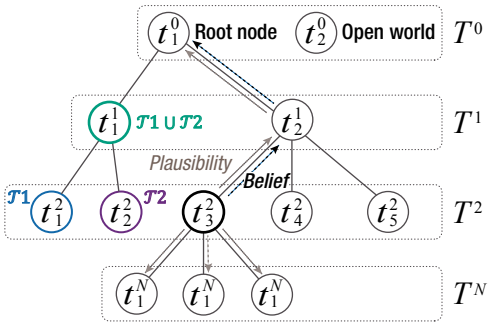

Figure 2. CID entanglement example. This example illustrates a situation resulting in CID entanglement between two crossing tracks. (Modified from Figure 1.[9])

the *BetP* distribution for both tracks represents a high confidence in the CID at this level.

When the tracks cross at time τ = 1, it is no longer clear which CID to attribute to which track. Historically, there have been two ways to address this challenge: (1) simply designate both tracks as having an unknown CID or (2) ignore the ambiguity and leave the CID estimates on each track unchanged. Both approaches are problematic. In the first case, the end user now has lost any knowledge of the CID estimates previously acquired. In the latter case, the CID estimate assigned to each track is 50% likely to be incorrect. Since a tracking ambiguity represents a case where at least one BBA is reliably associated with each track but we do not know which one, the disjunctive combination rule is used to compute the CID estimate of each track. Another way to justify the use of the disjunctive rule is to leverage set theory.[10] When tracks $\mathcal{T}1$ and $\mathcal{T}2$ cross (become ambiguous), we know that either BBA $m_{\mathcal{T}1(\tau=0)}$ OR $m_{\mathcal{T}2(\tau=0)}$ is correctly associated to each track, and therefore we apply disjunctive logic to combine the information represented by these BBAs. By applying the disjunctive combination rule to the BBAs $m_{\mathcal{T}1(\tau=0)}$ and $m_{\mathcal{T}2(\tau=0)}$, we obtain

	$m_{\mathcal{T}2(\tau=0)}\{t_1^2\} = .1$	$m_{\mathcal{T}2(\tau=0)}\{t_2^2\} = .9$
$m_{\mathcal{T}1(\tau=0)}\{t_1^2\} = .9$	$m_{(\mathcal{T}1\cup\mathcal{T}2)(\tau=1)}(\{t_1^2\}) = .9 * .1 = .09$	$m_{(\mathcal{T}1\cup\mathcal{T}2)(\tau=1)}(\{t_1^2, t_2^2\}) = .9 * .9 = .81$
$m_{\mathcal{T}1(\tau=0)}\{t_2^2\} = .1$	$m_{(\mathcal{T}1\cup\mathcal{T}2)(\tau=1)}(\{t_1^2, t_2^2\}) = .1 * .1 = .01$	$m_{(\mathcal{T}1\cup\mathcal{T}2)(\tau=1)}(\{t_2^2\}) = .9 * .1 = .09$

$$m_{(\mathcal{T}1\cup\mathcal{T}2)(\tau=1)}(\{t_1^2, t_2^2\}) = m_{(\mathcal{T}1\cup\mathcal{T}2)(\tau=1)}(\{t_1^1\}) = .81 + .01 = .82 \quad (9)$$

$$m_{(\mathcal{T}1\cup\mathcal{T}2)(\tau=1)}(\{t_2^2\}) = .09,$$

which represents the new disjunctively combined belief masses for both tracks $\mathcal{T}1$ and $\mathcal{T}2$, since we do not know which track to attribute the evidence to after their crossing. The *BetP* over the hypotheses in level T^2 is

$$BetP_{(\mathcal{T}1\cup\mathcal{T}2)(\tau=1)}(\{t_1^2\}) = .09 + \frac{1}{2}(.82) = .5$$
$$BetP_{(\mathcal{T}1\cup\mathcal{T}2)(\tau=1)}(\{t_2^2\}) = .09 + \frac{1}{2}(.82) = .5,$$

and the *BetP* over the hypotheses in level T^1 is

$$BetP_{(\mathcal{T}1\cup\mathcal{T}2)(\tau=1)}(\{t_1^1\} = \{t_1^2, t_2^2\}) = \frac{.82}{1-.18} = \frac{.82}{.82} = 1.$$

From these *BetP* calculations, it is clear that the CID estimate should be reported at level T^1 and the resulting estimate is t_1^1, which is the parent node of the CID estimates for each of the crossing tracks. Thus, in this situation, both tracks should be given the new CID estimate resulting from the disjunctive combination rule. Note that while the specificity of the CID estimate for each of the crossing tracks is diluted, it is not inconsistent with either of the original tracks and it does not represent a complete loss of knowledge by simply declaring that the CID is unknown. Thus, the general rule for CID estimation in cases of association ambiguity is to apply the TBM disjunctive combination rule (Eq. 3) and assign the result to all the ambiguous tracks.

Following the separation of the two tracks after their crossing at time τ = 1, consider the case where, at time τ = 2, $\mathcal{T}2$ is observed by a sensor and has new evidence unambiguously associated with it, which is represented by the following BBA:

$$\begin{matrix} m_{\mathcal{T}2(\tau=2)}\{t_1^2\} = .1 \\ m_{\mathcal{T}2(\tau=2)}\{t_2^2\} = .9 \end{matrix}. \quad (10)$$

The updated CID estimate, since we are assuming the sensor evidence is unambiguously associated with $\mathcal{T}2$, is found by taking the conjunctive combination of the sensor evidence and the combined BBA from Eq. 6. This leads to a combined BBA of

	$m_{\mathcal{T}2(\tau=2)}\{t_1^2\} = .1$	$m_{\mathcal{T}2(\tau=2)}\{t_2^2\} = .9$
$m_{(\mathcal{T}1\cup\mathcal{T}2)(\tau=1)}(\{t_1^2\}) = .09$	$m_{((\mathcal{T}1\cup\mathcal{T}2)\cap\mathcal{T}2)(\tau=2)}\{t_1^2\} = .009$	$m_{((\mathcal{T}1\cup\mathcal{T}2)\cap\mathcal{T}2)(\tau=2)}\{\emptyset\} = .081$
$m_{(\mathcal{T}1\cup\mathcal{T}2)(\tau=1)}(\{t_1^2, t_2^2\}) = m_{(\mathcal{T}1\cup\mathcal{T}2)(\tau=1)}(\{t_1^1\}) = .82$	$m_{((\mathcal{T}1\cup\mathcal{T}2)\cap\mathcal{T}2)(\tau=2)}\{t_1^2\} = .082$	$m_{((\mathcal{T}1\cup\mathcal{T}2)\cap\mathcal{T}2)(\tau=2)}\{t_2^2\} = .738$
$m_{(\mathcal{T}1\cup\mathcal{T}2)(\tau=1)}(\{t_2^2\}) = .09$	$m_{((\mathcal{T}1\cup\mathcal{T}2)\cap\mathcal{T}2)(\tau=2)}\{\emptyset\} = .009$	$m_{((\mathcal{T}1\cup\mathcal{T}2)\cap\mathcal{T}2)(\tau=2)}\{t_2^2\} = .081$

$$m_{((\mathcal{T}1\cup\mathcal{T}2)\cap\mathcal{T}2)(\tau=2)}\{t_1^2\} = .009 + .082 = .091$$
$$m_{((\mathcal{T}1\cup\mathcal{T}2)\cap\mathcal{T}2)(\tau=2)}\{t_2^2\} = .738 + .081 = .819$$
$$m_{((\mathcal{T}1\cup\mathcal{T}2)\cap\mathcal{T}2)(\tau=2)}\{\emptyset\} = .081 + .009 = .09$$

and a corresponding *BetP* at level T^2 of

$$BetP_{((\mathcal{T}1\cup\mathcal{T}2)\cap\mathcal{T}2)(\tau=2)}(\{t_1^2\}) = \frac{.091}{1-.09} = \frac{.091}{.91} = .1$$
$$BetP_{((\mathcal{T}1\cup\mathcal{T}2)\cap\mathcal{T}2)(\tau=2)}(\{t_2^2\}) = \frac{.819}{1-.09} = \frac{.819}{.91} = .9,$$

which is the same *BetP* distribution that $\mathcal{T}2$ had before the tracks' crossing at time $\tau = 0$ (Eq. 8). In this situation, the new evidence has almost completely disentangled the CID ambiguity introduced by crossing at time $\tau = 1$. While the *BetP* distribution at time $\tau = 2$ is the same as at time $\tau = 0$, the belief mass distribution is slightly different from the original. This is due to the BBA for $\mathcal{T}2$ (Eq. 10) having a belief mass of less than 1 for the hypothesis t_2^2. However, in general, new evidence incorporated after an ambiguity can disentangle some of the dilution caused by the ambiguity. Thus, the general rule for CID estimation where there is no ambiguity is to apply the TBM conjunctive combination rule (Eq. 2) and assign the result to the track for which the evidence was associated.

Efficient Implementation

The TBM combination rules fundamentally consist of set operations, specifically set intersection and union. This is particularly convenient when employing TBM using a taxonomy of hypotheses since a taxonomy concisely captures overlapping and nonoverlapping hypothesis sets. One efficient way to represent sets and perform set operations is to encode the sets using BitSets.[11] If a hypothesis taxonomy consists of N leaf nodes, then each hypothesis can be presented within the taxonomy as a BitSet of size N. An example of a simple BitSet encoding is presented in Figure 3 using a subset of the CoT classification taxonomy. The figure shows a portion of the CoT taxonomy where each node has a string defined by CoT (e.g., a–.–A) as well as the corresponding BitSet representation that is generated as follows. First, given a taxonomy with S leaf nodes, we specify that each BitSet representing each node in the taxonomy is of fixed length S. Starting with the leaf nodes, we define a BitSet with precisely one "true" bit while ensuring that no two BitSets contain the "true" bit in the same bit slot (i.e., the intersection between any two BitSets is the zero BitSet). Then, for each non-leaf node in the taxonomy, the node is represented by a BitSet that is equal to the union of all leaf-node BitSets contained as a subset in that node. As an example, the CoT string for the taxonomy node representing all "air" target types is a–.–A, while its BitSet representation is 1111000000 since this node represents four leaf nodes. Note also that the BitSet representation for a taxonomy is not unique as the leaf-node BitSets can be generated in many different ways and still satisfy the above-stated constraints.

Each leaf-node hypothesis (e.g., equipment, unit, structure) in the taxonomy has a BitSet representation consisting of one "1" and the remaining bits set to "0" where the position of the "1" is unique for each leaf node. Any non-leaf node in the taxonomy is the union of the BitSets of all its children nodes (e.g., air, ground). The root node, which represents the universe of hypotheses, is represented by a BitSet with all N values equal to "1." BitSet intersection and union operations are performed by comparing the values (0 or 1) in each position between two BitSet values. If there is any position where both BitSets have a "1," then the intersection BitSet contains a "1" in that position and the intersection is non-empty. If there is any position where either BitSet has a "1," then the union BitSet contains a "1" in that position and the union is non-empty. When evaluating the conjunctive combination rule, note that BitSet intersections that are non-empty represent non-zero terms in the summation of Eq. 2. When evaluating the disjunctive combination rule, note that BitSet unions that are non-empty represent non-zero terms in the summation of Eq. 3. As a result, the use of a taxonomy to represent the hypothesis space combined with BitSet representations of each node in the taxonomy makes the evaluation of TBM combination rules conceptually straightforward and computationally efficient.

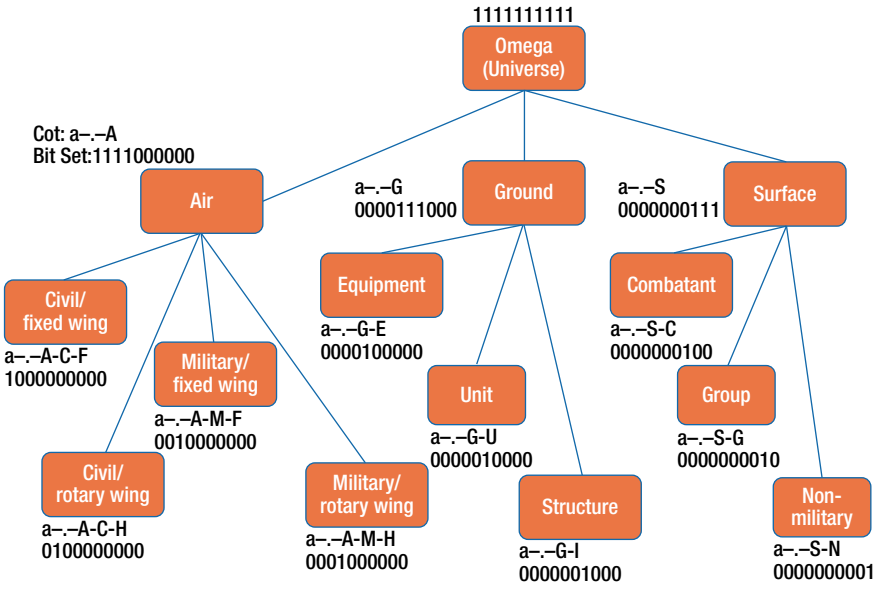

Figure 3. Using BitSets to represent hypotheses. This example of a simple BitSet encoding uses a subset of the CoT classification taxonomy, where each node has a string defined by CoT (e.g., a–.–A) as well as the corresponding BitSet representation that is generated.

Multi-INT Features and Evidence

In the above examples of TBM combination rules, we assumed that each contributing source provided a BBA over a set of nodes within the hypothesis space, but it did not reveal how the BBA was computed. To support a multi-INT solution, we need to relate the features for each INT to the hypothesis space. To this end, we introduce the concept of CID features and compute the CID evidence based on estimates of those features. Both categorical (discrete) and continuous features may be necessary to capture the BBA for a given INT source. No matter the source type, the estimated features need to be able to be mapped to the hypothesis space in order to impart evidence. For an electronic intelligence (ELINT) source, the features might be the electronic notation (ELNOT) of a signal. Because ELNOTs do not uniquely identify a battlespace object, an ELNOT estimate imparts evidence on many hypotheses in the taxonomy. Suppose, for example, that an estimated ELNOT identifies a military aircraft (both fixed-wing and rotary-wing) at a confidence level c. Under the example CoT taxonomy in Figure 3, the resultant BBA would be

$$\begin{aligned} m_{ELNOT}\{\{a-.-A-M-F, a-.-A-M-H\}\} \\ = m_{ELNOT}\{\{0010000000,0001000000\}\} \\ = m_{ELNOT}\{0011000000\} = c \end{aligned}.$$

$$m_{ELNOT}\{\{\emptyset\}\} = 1 - c$$

For a geospatial intelligence (GEOINT) source, the features might be the size or shape of an imaged battlespace object. In this case, the estimate (e.g., length) is a continuous variable, and the corresponding BBA is computed by the overlap between the values of the estimated feature and the known features of battlespace objects. As a further example using kinematics, consider a source that estimates the maximum speed of a battlespace object (see Figure 4), resulting in a mean and standard deviation (accounting for the inherent uncertainty in the ability to observe speed). This can be captured as a Gaussian distribution over the speed feature. The possible speeds of various vehicles are known and can be captured using intervals over the speed feature space. Then, the evidence (i.e., BBA) is given by evaluating the overlap of the Gaussian with the intervals of speed.

In the example shown in Figure 4, the BBA for the max speed feature over the CoT taxonomy is

$$\begin{aligned} m_{ELNOT}\{\{a-.-A-C-F, a-.-A-M-F\}\} = .9 \\ m_{ELNOT}\{\{a-.-A-W-M\}\} = .1 \end{aligned}.$$

That is, ~90% of the Gaussian distribution overlaps the region corresponding to fixed-wing aircraft, while ~10% of the Gaussian distribution overlaps the region corresponding to a missile, and there is negligible overlap with the other regions.

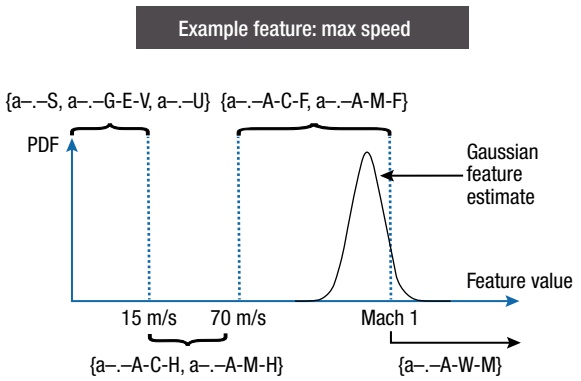

Figure 4. Example showing max speed continuous feature and evidence assignment. This example using kinematics illustrates a source that estimates the maximum speed of a battlespace object, resulting in a mean and standard deviation (accounting for the inherent uncertainty in the ability to observe speed).

Adaptive Output

The last remaining requirement for CID estimation is the ability to produce an actionable output that is both informative and concise. Once the TBM is used to combine evidence and compute the BBA over the taxonomy of hypotheses, the PPT (*BetP*) can be computed over any collection of exhaustive and mutually exclusive hypotheses (i.e., a reporting hypothesis space) within the taxonomy. As a matter of procedure, the PPT (Eq. 5) is applied over multiple reporting hypothesis spaces within the taxonomy and a decision is made as to which hypothesis space to use as the reported result. Since the objective of the adaptive output is to provide the most concise output without significant loss of information, an information theoretic approach derived from the domain of data compression is applied.[12] The procedure is as follows:

- The probability distribution (*BetP*) is computed across the set of leaf nodes (R_0) within the taxonomy, which serves as the reference probability distribution since it is the most specific representation of the FOD. The *BetP* formula for an arbitrary reporting hypothesis space R_i is

$$BetP^{R_i}(A) = \sum_{B \in 2^{R_i}} \frac{|A \cap B|}{|B|} \frac{m^{R_i}(B)}{1 - m^{R_i}(\emptyset)} \; \forall A \in R_i, i \geq 0.$$

- The Kullback–Leibler (KL) divergence,[13] $D(R_i||R_0)$, between a reporting hypothesis space (R_i) and the leaf node (reference) hypothesis space (R_0) in the taxonomy is computed by writing the KL divergence in terms of the Shannon entropy of the reference hypothesis space, $H(R_0)$, and the cross entropy between the reporting hypothesis space and the reference hypothesis space, $H(R_i,R_0)$, as follows.[13]

  Let $N(B)$ be the number of leaf nodes that represent hypothesis B in a reporting space R_i with $i > 0$, then:

$$H(R_0) = -\sum_{A \in R_0} BetP^{R_0}(A)\, ln\big(BetP^{R_0}(A)\big)$$

$$H(R_i, R_0) = -\sum_{A \in R_0; B \in R_i; A \cap B = A} BetP^{R_0}(A)\, ln\big(P^{R_i}(B)\big)$$

with

$$P^{R_i}(B) = \frac{BetP^{R_i}(B)}{N(B)} \quad \forall B \in R_i$$

and

$$D(R_i||R_0) = H(R_i, R_0) - H(R_0).$$

- Report the probability distribution $BetP^{R_i}$ for hypothesis space R_i that is the most concise (i.e., has the least number of non-zero probabilities) and satisfies the specified type II error, α, where the type II error is given by Stein's lemma:[13]

$$P^{R_i}_{Type\ II} = 1 - exp\{-D(R_i||R_0)\} < \alpha.$$

Generally, the *BetP* distribution over a smaller reporting hypothesis space will contain less information than the *BetP* distribution over a larger reporting hypothesis space. That is, reporting a CID estimate higher up in the taxonomy generally leads to information loss when compared to the leaf nodes. However, reporting the *BetP* distribution over the leaf nodes may not be concise. The type II error, α, allows for reporting at a more concise level in the taxonomy with a controllable amount of information loss. In this context, the type II error probability, α, is the probability of deciding to report on hypothesis space R_i when the data supports selecting the reference hypothesis space R_0 because R_i is too concise to represent the same information as R_0.

We applied this procedure to a CoT taxonomy consisting of 754 leaf nodes, using real evidence from three features: (1) maximum speed, (2) maximum altitude, and (3) friend or foe declaration. We defined an allowable type II error of 0.05 (5% compression loss). The evidence resulted in 19 (of 754) leaf nodes having non-zero *BetP* probability and a resultant entropy, $H(R_0)$, of 2.94. Reporting 19 CID estimates and the corresponding *BetP* distribution to a warfighter may be overwhelming. The next level higher in the taxonomy has 34 nodes and a *BetP* distribution with 2 non-zero values. The cross-entropy, $H(R_1,R_0)$, is also identically 2.94; thus, the KL divergence, $D(R_1||R_0)$, and the type II error, $P^{R_1}_{Type\ II}$, is 0. As a result, the CID estimate can be reported as two hypotheses with the corresponding *BetP* probabilities without any information loss relative to reporting the *BetP* for the leaf nodes. This occurs when all the leaf nodes having non-zero probability are equivalently represented by a single parent node within the taxonomy (i.e., sibling leaf nodes with non-zero probability have equal probability). To complete the example, the next level higher in the taxonomy has five hypotheses and a *BetP* distribution with one non-zero value. However, the cross-entropy, $H(R_2,R_0)$, is 4.04; thus, the KL divergence, $D(R_2||R_0)$, is 1.1 and the type II error, $P^{R_2}_{Type\ II}$, is 0.4. Since this exceeds the specified type II error level, the process is terminated and the CID estimate *BetP* distribution of R_1 is reported to the warfighter.

DEALING WITH LATENT CID INFORMATION

Real-world targeting applications often involve situations where highly informative CID information can arrive at the decision-maker with significant latency compared to other tracking information. This can occur for a variety of reasons, including processing and communications latencies, depending on the sensing phenomenology. For the warfighter to exploit this latent CID information, two algorithm challenges must be addressed. First, an algorithm for associating the latent information with historical tracking data is required. Second, the CID estimate must account for possible tracking ambiguities between the time of association and the decision time of the warfighter. Latent data association requires maintenance of sufficient time history of tracking data, while CID estimation at a later time requires an accounting of periods of ambiguity. To address both challenges, a track segment graph (TSG)[14] is used to retain a compressed representation of the movement of each tracked object, including ambiguities, and the time periods associated with each node in the TSG are used to form an interval binary search tree (IBST).[15] Given latent CID information (and the associated "valid" time of that information), the IBST is searched to determine all nodes in the TSG that existed at the valid time, and the TSG is used to determine whether those nodes represent periods of tracking ambiguity. If an unambiguous association is made with a node containing a single track, the TSG is traversed from the association time to the current time and a CID estimate is computed using the TBM algorithms detailed in a previous section (CID Estimation Rules) based on any ambiguities that may have occurred.

For the purposes of this article, we define an unambiguous association as an association between evidence and tracks that has a sufficiently low probability of being incorrect. For example, if receiving an image of a battlespace object, we can say that it is unambiguously associated with a track if the geographic location of the image and a track, accounting for their uncertainties, are such that the image is highly unlikely to be geospatially consistent with another track. Alternatively, an ambiguous association is declared when there is sufficiently high likelihood of geospatial consistency between an image and more than one track. While they are not the emphasis of this article, tests such as the Mahalanobis distance[16] between two Gaussian estimates (one

representing the geographic location of CID evidence and the other representing the geographic location of a track) can be used with appropriate type I and type II error thresholds to determine whether an association is ambiguous or unambiguous. In this context, a type I error is the declaration of no association when the CID evidence originated from the tracked object (missed association), and a type II error is a declaration of association when the CID evidence did not originate from the tracked object (false association).

A TSG is created iteratively as tracks are updated in order to capture linkages between tracks that become ambiguous (i.e., closely spaced and moving in the same direction) at different points in time. Here, similar to the definition of association ambiguity between CID evidence and tracks above, we consider two (or more) tracks to be ambiguous if their geographic location and velocity estimates (along with their uncertainties) are similar enough that there is a high probability of these tracks representing the same battlespace object. An ambiguity is declared when the Mahalanobis distance is smaller than a threshold value derived from a specified hypothesis testing error, β. The smaller the value of β, the more similar the track state estimates must be to declare an ambiguity. If comparing tracks that have six-state estimates (3-D position and 3-D velocity) with a desired probability of incorrectly declaring an ambiguity to be less than 0.1, then β is 10.645, given by a chi-square distribution[17] with six degrees of freedom and a critical value of 0.9. Thus, when two or more tracks have a Mahalanobis distance of less than 10.645, they will be declared ambiguous and an ambiguous node will be added to the TSG with the ambiguous tracks. Additionally, the Mahalanobis test can be used to determine when an ambiguity is resolved and tracks can be placed in their own unambiguous node in the TSG. Here it is useful to build in some hysteresis to avoid toggling between ambiguous and unambiguous nodes in the TSG. For example, the chi-square threshold for resolving an ambiguity may use a critical value of 0.95, which is 12.592. Thus, once tracks have a Mahalanobis distance exceeding 12.592, they can be declared unambiguous and will be placed in an unambiguous node in the TSG corresponding to a single track.

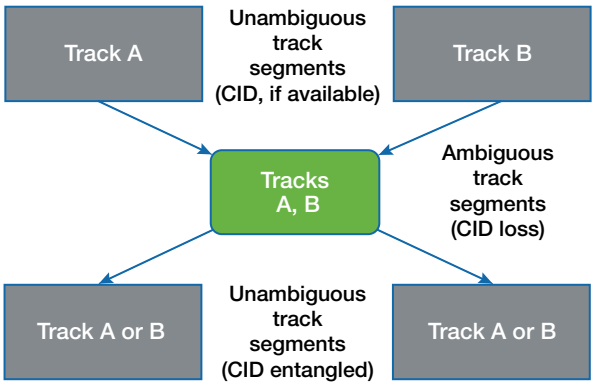

Figure 5. Example TSG illustrating crossing tracks. At the top of the graph are two nodes representing the time periods for which tracks A and B are well separated (i.e., unambiguous). At the bottom of the graph, the tracks have separated spatially and are no longer ambiguous; however, the CID information is now entangled since it is not clear which outgoing track should have which initial CID information attributed to it.

The above-described approach to using the Mahalanobis distance to build a TSG has been used to reduce computational burden for multi-hypothesis tracking algorithms that evaluate hypotheses over extended time horizons.[14] Figure 5 presents a TSG for the case of two crossing tracks as discussed in detail in the section on CID estimation rules. At the top of the figure are two nodes representing the time periods for which tracks A and B are well separated (i.e., unambiguous). The CID estimate for tracks A and B at this time will be computed using the conjunctive combination rule (Eq. 2) on all BBAs derived from CID evidence (see the Multi-INT Features and Evidence section) associated with these tracks. When the tracks become too close according to a Mahalanobis distance test, they are declared to be ambiguous and a single node is formed, which represents both tracks A and B. If tracks A and B had CID information, the ambiguity indicates loss of the ability to attribute this information to either track uniquely, and therefore both tracks A and B would be assigned a CID estimate using the disjunctive combination rule (Eq. 3) across all BBAs previously associated with both tracks. Finally, at the bottom of the graph, the tracks have separated spatially and are no longer ambiguous; however, the CID information is now entangled since it is not clear which outgoing track should have which initial CID information attributed to it. That is, both tracks will still retain a CID estimate resulting from the disjunction performed when they became ambiguous.

Each node within the TSG generally contains the following information:

- Time period represented by the node
- Track identifiers (e.g., track numbers)
- CID evidence (e.g., BBAs over a reference taxonomy)
- Track state representation (e.g., set of track states, polynomial curve fit)

Using the TSG can greatly reduce computational and storage requirements when compared to simply storing all track updates.[14] Computational savings are achieved by determining ambiguities as track updates occur instead of having to reassess the track picture for ambiguities each time latent data are received, which could result in significant duplicate processing for moderate to large tracking problems. In addition, storage requirements can

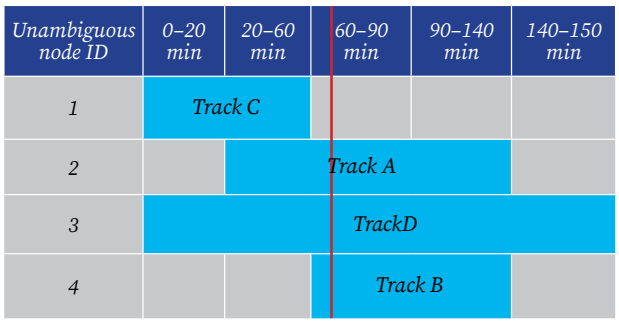

Figure 6. Timeline of unambiguous track history graph nodes over 150 min of history. Time periods in blue indicate unambiguous nodes, each representing a single track within the TSG. The red line indicates the point at which the IBST is queried for all unambiguous nodes.

be reduced by representing each node in the TSG by a polynomial curve fit or some other compressed representation of the track state over that time period. For example, if a track is moving in a straight line, the node may simply contain the parameters to describe this line rather than many track states that represent the time series along this line. With the computational and storage requirements reduced, latent data can be associated in complex tracking environments and with low computational latency.

While the TSG is being generated, a corresponding IBST is maintained to store the time intervals over which all unambiguous nodes in the TSG exist. We focus on unambiguous nodes because we cannot uniquely attribute CID evidence with ambiguous tracks. When latent data are received, the IBST allows for quick identification of unambiguous nodes that temporally intersect with the timestamp on the latent data. The IBST provides an $O(log(N))$ complexity search[18] for candidate unambiguous nodes to associate with latent CID evidence. Figure 6 provides an example IBST that has been generated over the course of 150 min, where the time periods in blue indicate unambiguous nodes, each representing a single track within the TSG. So tracks represented by node 1 (track C), node 2 (track A), node 3 (track D), and node 4 (track B) in the TSG were unambiguous with all other tracks between 0 and 60 min, 20 and 140 min, 0 and 150 min, and 60 and 140 min, respectively. Also, assume that nodes 2 and 4 are merged at the 140-min mark due to ambiguity between tracks A and B. This is the same scenario as depicted in Figure 5.

At 150 min, a latent geospatial detection with CID evidence (e.g., an image chip with location information) collected at 70 min is received by the CID estimation process. We query the IBST for all unambiguous nodes at time 70 min (red line in Figure 6), and the IBST returns nodes 2, 3, and 4 as candidates for association with the CID evidence. A final association of CID evidence is only made if an unambiguous association can be made with one of the tracks represented by one of these nodes. Again, this decision can be made using the Mahalanobis distance test.

Furthering the example, assume an unambiguous association of CID evidence is made with the track represented by node 2 (track A). The associated CID evidence is combined with any other CID evidence for track A using the TBM conjunctive combination rule (Eq. 2). To update track A's CID to the current time (150 min), the TSG is traversed (e.g., from top to bottom in Figure 5) starting from node 2 forward in time until the current time. According to Figure 6, node 2 is no longer unambiguous (i.e., it has been merged with track B to form an ambiguous node) at time 140 min, and therefore, all CID evidence associated with node 2 (track A) is fused with all CID evidence attributed to node 4 (track B) using the TBM disjunctive combination rule (Eq. 3).

Figure 7 depicts the process of CID estimation after the latent CID evidence has been associated with track A at 70 min. Before the latent association in our example, track A had only two features (e.g., max altitude and max speed) attributed to it. After the latent association, track A now has three features (e.g., max altitude, max speed, and length extracted from an image chip) attributed to it (feature 3 highlighted in blue). Thus, we start by conjunctively combining the BBAs for the evidence ($m_{A.1}$, $m_{A.2}$, and $m_{A.3}$) from these features (see the Multi-INT Features and Evidence section) to compute the belief mass m_A for track A. Continuing to traverse the TSG, we find an ambiguous node at time 140 min that consists of tracks A and B. So, in accordance with the section on CID estimation rules, we now conjunctively combine the CID evidence attributed to track B ($m_{B.1}$ and $m_{B.2}$), resulting in m_B, and then the belief masses for tracks A and B will be disjunctively

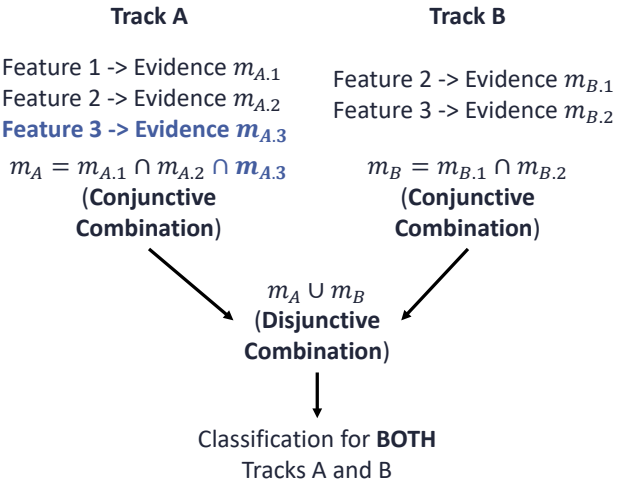

Figure 7. Evidence combination using TBM for an ambiguous node containing two tracks. This figure depicts the process of CID estimation after the latent CID evidence has been associated with track A at 70 min.

combined, $m_A \cup m_B$, to produce the CID estimation for both track A and track B. An interesting result of this example is that the CID estimate of track B is affected by the latent association of CID evidence to track A through the ambiguity of tracks A and B.

If, after 150 min, these tracks become unambiguous again (such as in Figure 5), the CID estimate of both tracks will remain equal to $m_A \cup m_B$ as in Figure 7 because we are still unable to unambiguously attribute the evidence BBAs $m_{A.1}$, $m_{A.2}$, $m_{A.3}$, $m_{B.1}$, and $m_{B.2}$ to the correct track. If, however, CID features are associated with track A or B after their ambiguity is resolved, then the corresponding CID evidence BBA is conjunctively combined with the current CID evidence for each track, $m_A \cup m_B$.

CONCLUSIONS

This article presented an automated CID estimation process that accommodates imprecise CID evidence and heterogeneous multi-INT CID feature spaces to generate an actionable result suitable for targeting decisions. The TBM provides a framework for combining uncertain, ambiguously associated, and conflicting evidence while accepting an incomplete hypothesis space (open-world assumption). The TBM was applied to a taxonomy with hierarchical hypothesis spaces, allowing for an efficient implementation using BitSet encoding. Given the inherent variation in output specificity allowed by a taxonomy, this article presented an information theoretic approach to selecting the most concise hypothesis space while adhering to suitable compression loss. Finally, a TSG and IBST representation of historical track relationships was used to address the association of latent CID evidence and CID estimation in the presence of subsequent tracking ambiguities. Together, these algorithms provide a robust decision aid for targeting that overcomes historical deficiencies that have hindered the real-world use of automated CID estimation.

REFERENCES

[1]"Tactical Data Link (TDL) 16 message standard," MIL-STD-6016, Department of Defense, Washington, DC, 2002, https://quicksearch.dla.mil/qsDocDetails.aspx?ident_number=123964.

[2]"Joint military symbology," MIL-STD-2525D, Department of Defense, Washington, DC, 2014, https://www.jcs.mil/Portals/36/Documents/Doctrine/Other_Pubs/ms_2525d.pdf.

[3]M. Butler, "The developer's guide to Cursor on Target," MITRE Corp., Bedford, MA, 2005, https://apps.dtic.mil/sti/citations/ADA637348.

[4]J. H. Kim and J. Pearl, "A computational model for causal and diagnostic reasoning in inference systems," in *Proc. 8th Int. Joint Conf. Artif. Intell.*, 1983, vol. 1, pp. 190–193, https://dl.acm.org/doi/10.5555/1623373.1623417.

[5]A. P. Dempster, "A generalization of Bayesian inference," *J. Roy. Statist. Soc.: Ser. B*, vol. 30, no. 2, pp. 205–232, 1968, https://doi.org/10.1111/j.2517-6161.1968.tb00722.x.

[6]P. Smets, "Analyzing the combination of conflicting belief functions," *Inf. Fusion*, vol. 8, no. 4, pp. 387–412, 2007, https://doi.org/10.1016/j.inffus.2006.04.003.

[7]P. Smets, "The combination of evidence in the transferable belief model," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol., 12, no. 5, pp. 447–458, 1990, https://doi.org/10.1109/34.55104.

[8]P. Smets and R. Kennes, "The transferable belief model," *Artif. Intell.*, vol. 66, no. 2, pp. 191–234, 1994, https://doi.org/10.1016/0004-3702(94)90026-4.

[9]W. Farrell III and A. M. Knapp, "Multisource taxonomy-based classification using the transferable belief model," in *Proc. SPIE 8407, Multisensor, Multisource Inf. Fusion: Architectures, Algorithms, Appl. 2012*, 840704, https://doi.org/10.1117/12.923873.

[10]J. Bagaria, "Set theory," in *Stanford Encyclopedia of Philosophy*, spring 2020 ed., Edward N. Zalta, Ed., https://plato.stanford.edu/archives/spr2020/entries/set-theory/.

[11]A. Wellings, "Multiprocessors and the real-time specification for Java," in *Proc. 2008 11th IEEE Int. Symp. Object Compon.-Oriented Real-Time Distrib. Comput. (ISORC)*, pp. 255–261, https://doi.org/10.1109/ISORC.2008.22.

[12]A. Kaltchenko, "Reexamination of quantum data compression and relative entropy," *Phys. Rev. A*, vol 78, no. 2, article 22311, 2008, https://doi.org/10.1103/PhysRevA.78.022311.

[13]S. Kullback, *Information Theory and Statistics*. New York: Dover Publications, 1997.

[14]C.-Y. Chong, G. Castanon, N. Cooprider, S. Mori, R. Ravichandran, and R. Macior, "Efficient multiple hypothesis tracking by track segment graph," in *Proc. 2009 12th Int. Conf. Inf. Fusion*, 2009, pp. 2177–2184, http://fusion.isif.org/proceedings/fusion09CD/data/papers/0517.pdf?

[15]T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*, 3rd ed. Cambridge, MA: MIT Press and McGraw-Hill, 2009.

[16]R. De Maesschalck, D. Jouan-Rimbaud, and D. L. Massart, "The Mahalanobis distance," *Chemometrics Intell. Lab. Syst.*, vol. 50, no. 1, pp. 1–18, 2000, https://doi.org/10.1016/S0169-7439(99)00047-7.

[17]A. Mood, F. A. Graybill, and Duane C. Boes, *Introduction to the Theory of Statistics*, 3rd ed., pp. 241–246. Noida, India: McGraw-Hill, 1974.

[18]D. E. Knuth, "Optimum binary search trees," *Acta Informatica*, vol. 1, no. 1, pp. 14–25, 1971, https://doi.org/10.1007/BF00264289.

William J. Farrell III, Force Projection Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

William J. (Jim) Farrell is a lead engineer for intelligence, surveillance, and reconnaissance (ISR) sensors and data fusion in APL's Force Projection Sector. He has a BS in physics and math from Towson University and an MS in physics from the University of Maryland, College Park. Jim has 20 years of experience in cutting-edge research and development of tracking, classification, and information fusion algorithms in space, air, land, surface, and subsurface warfare domains. He has been the principal investigator or technical lead for several Department of Defense programs as well as APL independent research and development programs, including serving as lead developer for fusion systems operating at four land sites on three continents and onboard dozens of Navy ships. A recognized information fusion subject-matter expert, he is the author of nearly 20 publications related to information fusion. He is a senior member of IEEE and served on the Lockheed Martin Joint Strike Fighter (F-35) corporate-wide tiger team for data fusion and the Army Research Laboratory's fusion advisory panel. Jim, along with teammates, won a 2020 APL Mission Accomplishment Award (Current Challenge) for work on Maritime Targeting Cell – Afloat (MTC-A), an initiative to enhance the US Navy's ability to provide long-range targeting. His email address is jim.farrell.III@jhuapl.edu.