

AD-A256 123

TION PAGE

Form Approved
OPM No. 0704-0188Public report
maintaining
for reducing
the Office ofper response, including the time for reviewing instructions, searching existing data sources gathering and
its regarding this burden estimate or any other aspect of this collection of information, including suggestions
tion Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to
Washington, DC 20503.

1. AGENCY USE ONLY (Leave Blank)	2. REPORT DATE 1991	3. REPORT TYPE AND DATES COVERED	
4. TITLE AND SUBTITLE Using An Assumption-Based Truth Maintenance System to Switch Context during Data Fusion Processing		5. FUNDING NUMBERS	
6. AUTHOR(S) Cronin, T.		7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) U.S. Army CECOM Signals Warfare Directorate Vint Hill Farms Station Warrenton, VA 22186-5100	
8. PERFORMING ORGANIZATION REPORT NUMBER 92-TRF-0040		9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	
10. SPONSORING/MONITORING AGENCY REPORT NUMBER		1. SUPPLEMENTARY NOTES	
12a. DISTRIBUTION/AVAILABILITY STATEMENT Statement A: Approved for public release; distribution unlimited.			
12b. DISTRIBUTION CODE			
13. ABSTRACT (Maximum 200 words) This report is organized into seven sections. The first deals with issues raised by the effect of time on discrete event modeling, and is known as the frame problem. The second section describes a classical control theory approach to multiple target tracking, which relies exclusively upon Bayes theorem and a Kalman filter to track objects which change spatial position over time. The third section on justification- based truth maintenance systems introduces the concept of retraction to correct temporal inconsistencies which arise in situation databases when objects change spatial position. The fourth section on assumption-based truth maintenance system is reactionary to the third, in that it proposes to relieve the computational burden of retraction by pursuing several feasible hypotheses in parallel. The fifth section, the problem of uncertainty in data fusion and the implication for logic-seeking truth maintenance technology is discussed. The sixth section describes some Army work which can provide real-time collateral inputs to Reid's target tracker, and also to the two styles of truth mainten- ance philosophy. The final section proposes a handshaking protocol to interface Reid's statistical system to a logic-driven truth maintenance system.			
14. SUBJECT TERMS Artificial Intelligence, Data Fusion, Frame Problem, Non-Monotonic Inference, Multiple Target Tracking Algorithm, Filters and Linear Tracking Models, Truth Maintenance Systems.		15. NUMBER OF PAGES 23	
16. PRICE CODE		17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	
18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED		19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	
20. LIMITATION OF ABSTRACT UL		21. LIMITATION OF ABSTRACT UL	

Using an Assumption-based Truth Maintenance System to Switch Contexts during Data Fusion Processing

T. Cronin
Signals Warfare Directorate
Warrenton VA 22186-5100

Accession For	
NTIS GR&I	☒
DTIC TAB	☐
Unannounced	☐
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

ABSTRACT

A contextual objective of data fusion is to monitor the temporal emergence and decay of battlefield events, to provide feedback to statistically-based correlation processes on the front end. It is suggested that de Kleer's assumption-based truth maintenance system (ATMS) is a natural technology for implementation of this task. An ATMS is well-suited to any task which entails diagnostic reasoning, and bears the distinction of being able to contrast premises (hard facts) with assumptions (tentative beliefs). However, its primary strength is its adeptness at maintaining multiple hypotheses in parallel. A battlefield event may be modeled as a function of the logical states of a set of Boolean switches which are either turned on or off whenever the event occurs. In practice, an event may be differentiated from other events by diagnosing spatial and temporal functions of sensor measurements arriving from the perceptual-level front end of the reasoning system. However, to implement this approach, the frame problem must be met head-on. Although correlation of sensor-derived feature data has traditionally been performed with statistical techniques which are subject to type 1 (failure to fuse) and type 2 (unwarranted fusion) errors, there is no technology available to undo the errors, should some inconsistency in the database arise as a result of adopting the result of the statistical test. For example, an inconsistency may arise because some previously stationary object becomes a mover, producing a feature vector which so significantly deviates from the mean that a failure to update the covariance matrix occurs, producing a type 1 statistical error. Whatever the cause, an inconsistency tends to compound error over time, thereby thwarting further analysis until it is resolved. If the act of fusing or failing to fuse two error ellipse probables is viewed as an instance of what de Kleer calls an assumption, then a tightly-controlled ATMS may later modify the act, by backtracking to the point in time at which the statistical error occurred. The problematic null hypothesis decision may then be retracted to ensure consistency over the situation database. Subsequently, the logical argument (justification) produced by utilizing the obsolete fusion product may be undone and recomputed using forward-chaining on the new hypothesis. The net result is that statistical testing in the traditional sense is enhanced with a temporal capability to alter decision (switch contexts). It should be pointed out that an ATMS in general tries to avoid retraction, but when faced with limited resources must begrudgingly resort to it. To compare this approach to others which have been utilized traditionally for multiple hypothesis management, the ATMS is compared and contrasted with D.B. Reid's multiple target tracker published in the control theory literature. It is shown that both techniques might benefit from some new computational geometry results which have been recently developed, and that the two techniques acting in collusion may interact to produce a more powerful multiple target tracking system than either technique acting alone can provide.

92-24680



424856

13

92 9 03 076

BACKGROUND

In dynamic data fusion domains, there is a requirement for an observer to both maintain cognizance and report upon the spatial positions and postures of mobile objects within his area of interest, especially in the context of other mobile objects which exhibit locomotion against a background of some relatively fixed objects, which may collectively either impede, accelerate, or have negligible effect upon the ambulatory progress of the movers. Furthermore, for the intelligence collection process, the observer is constrained by the resolution of remotely positioned sensors at his disposal, which by their very nature are constrained to be measurement processes, subject to both range (depth of measurement) and azimuth (breadth of measurement) errors. At any particular instant, the observer must have the wherewithal to vocalize his particular set of *beliefs*, which in data fusion domains generally consists of a verbal description of the perceived current locations and postures of objects within the range of his sensors, as well as objects which may lie just beyond his region of interest.

Outline of the Paper

The paper is organized into seven sections. The first deals with issues raised by the effect of time on discrete event modeling, and is known as the frame problem. The second section describes a classical control theory approach to multiple target tracking, which relies exclusively upon Bayes theorem and a Kalman filter to track objects which change spatial position over time. The third section on justification-based truth maintenance systems introduces the concept of retraction to correct temporal inconsistencies which arise in situation databases when objects change spatial position. The fourth section on assumption-based truth maintenance systems is reactionary to the third, in that it proposes to relieve the computational burden of retraction by pursuing several feasible hypotheses in parallel. In the fifth section, the problem of uncertainty in data fusion and the implication for logic-seeking truth maintenance technology is discussed. The sixth section describes some Army work which can provide real-time collateral inputs to Reid's target tracker, and also to the two styles of truth maintenance philosophy. The final section proposes a handshaking protocol to interface Reid's statistical system to a logic-driven truth maintenance system. It is hoped that the combined package might produce a system adept at switching contexts, by exploiting whichever inputs, whether they be statistical or logical, present themselves to a data fusion analyst.

THE FRAME PROBLEM AND NON-MONOTONIC INFERENCE

The Frame Problem

Real world events are in general continuous phenomena, exhibiting smooth gradual changes as dynamic processes evolve over time. At certain moments, the constituent functional components which comprise an event may become spatially static. It is during just such moments, when a dynamically transpiring event exhibits a moment of spatial stability, that it may be possible to characterize the event in terms of a classical computer model. At times in between, when the constituent phenomena are in a more chaotic state due to logistic and operational considerations dependent upon the spatial environment, the modeling issue may well be undecidable, due to prohibitive combinatorial constraints.

If a battlefield event is represented with a set of Boolean switches (or slots) which are respectively either on or off based on whether the logical assertions which model the event are true or false, then such a model may be construed to be a frame, the indicators of which are either

supported by evidence or not. If an indicator slot for an event is supported by evidence, then that particular event slot is said to be *instantiated*. It is instantiated because an instance (specific value) of a generalized random variable has been discovered empirically, which triggers the corresponding event slot to be true. The slot's value may exhibit a spectrum of truth, from *absolute* (if the variable is logical), to *probabilistic* (if the variable assumes probabilistic values).

The term "frame problem" (ref. 12) was coined by Patrick Hayes, who is a world-renowned logician, noted in the AI world for his work in common-sense reasoning. Although the reference cited is now nearly twenty years old, the frame problem remains unsolved. The problem arises because with discrete event modeling of dynamic situations, we do not know precisely how to chain static snapshots of events together, nor do we know what happens spatially and temporally between snapshots. The process of event modeling cannot capture the temporal continuum between frames, when it is required to revoke belief in a once credible event frame, because another one subsequently does a better job of explaining the current set of data. In data fusion applications for the military, the frame problem is intimately related to multiple target tracking, in which several distinct objects behave individually to comply with some global strategy envisioned by policymakers. During multiple target tracking, it is legitimate to rhetorically ask the question "how are targets behaving during the time intervals in which sensors are not passively or actively receiving information?"

As an example of the frame problem, consider the Chancellorsville campaign during the Civil War. On May day in 1863, Robert E. Lee found himself and the Confederate Army in a desperate situation. He had already divided his Army by leaving General Early in Fredericksburg with half the Confederate troops. In the interim, Joseph Hooker had crossed the Rappahannock river with 70,000 Union troops to threaten Lee and Stonewall Jackson, who had only 40,000 troops at their disposal. On the evening of the first, Jackson proposed to Lee to once again divide the Confederates by marching Jackson's entire Corps thirteen miles around the Union's right side. The plan was implemented the following morning, and by the evening of the second of May, Jackson was in position, and began to assault the Union flank. During this whole operation, the only indicators available to Hooker that he was being flanked were a confusing encounter with Jackson's rear elements marching westwards at midday, and comments later in the afternoon by Union soldiers on the right that troop movements could be heard to the west. Both indicators were dismissed by Hooker, and as a result his Army was severely routed and forced back across the river. Unfortunately, Jackson, who had implemented one of the most brilliant flanking maneuvers in military history, was fatally wounded by his own men that very evening.

Non-monotonic Inference

Closely related to the frame problem is *non-monotonic inference*. Non-monotonic inference is perhaps most easily grasped by describing its negation: monotonic inference. Suppose that one is able to deduce a conclusion given a set of assertions. If one can continue to deduce the same conclusion no matter how many more assertions are added to the original set, then the underlying logical regime is termed *monotonic* inference. If on the other hand, at some point the conclusion may no longer be derived from the assertions, the logical regime is termed *non-monotonic* inference. At the root of non-monotonic reasoning is the fact that in dynamic worlds, the spatial constituent components which comprise an event may move about over a period of time. This means that deductions based on the particular whereabouts of an object may become invalid as the clock ticks between the discrete times during which the objects are momentarily glimpsed by sensors.

Prior to formal recognition of the frame problem in 1973, there was no provision for non-monotonic inference in an automated reasoning system. All data entering the system were expected to correspond to true assertions, so that any conclusion derived previously would thereafter remain true. Clearly, this naive style of evidence accrual was inadequate for complex real world problems, particularly those involving the dynamic measurement processes encountered in data fusion domains. The graphic at Figure 1 portrays both the simplicity and naivety of monotonic logic. Whenever an assertion is perceived to be true, it is made a permanent part of the database, and is true thereafter. Before the frame problem was formalized, this style of accruing evidence was prevalent in AI reasoning. Returning to our example from the Civil War, it is the same logic used by Joseph Hooker when he continued to believe Stonewall Jackson was with Lee to the South, when in fact Jackson had swung wide to the west with a flanking maneuver.




Figure 1. Naive treatment of belief accumulation.

REID'S MULTIPLE TARGET TRACKING ALGORITHM

Leaving the frame problem behind for a moment, let us turn to a traditional implementation of multiple target tracking. Multiple target tracking is one aspect of data fusion which is well-studied, as evidenced by a large number of citations in the literature. Previous work in the area has been primarily statistical in nature. Highly representative of the work is a seminal paper by D.B. Reid, which describes a Bayesian tracking algorithm (ref. 18). This is a powerful algorithm, in that it uses a Kalman filter (ref. 14) to predict new locations of objects, and uses a variant of Bayes decision rule to decide which datasets belong with what tracks. It is also noteworthy in that it is one of the first systems to provide for maintaining multiple hypotheses in parallel. When targets behave in a linear fashion, the technique is unparalleled in its predictive capability.

The algorithm makes use of sensor-derived measurements to develop the probability that a sample value corresponds to a known target, a false target, or a new target. The tracking logic is based on a linear recursive Kalman filter. Sensors are differentiated into two types, the first being an active sensor which seeks to either confirm or deny the existence of a set of targets, and which can provide a corresponding estimate of target distribution, and the second a passive sensor which is capable of producing a positive report of a single target only. Reid assumes that the previously known targets are binomially distributed, and that new and false targets are each Poisson distributed. His chief contribution is the following formula, which calculates the probability that a particular dataset is associated with a particular target:

$$P_i^k = \frac{1}{c} P_D^{N_{DT}} (1 - P_D)^{(N_{TOT} - N_{DT})} \beta_{FT}^{N_{FT}} \beta_{NT}^{N_{NT}} x \left[\prod_{m=1}^{N_{DT}} N(Z_m - H\bar{x}, \beta) \right] P_i^{k-1} \quad (1)$$

Although Reid's tracker is quite powerful for objects which adhere to a linear model, there are situations, especially for the ground-based Army, in which the technique falls short in its predictive capability. Reid enumerated six issues not to be addressed in his paper. The list consists of the following: a. nonlinear measurements; b. nonlinear dynamics; c. maneuvering targets; d. requirement for an adaptive algorithm; e. problems of multiple sensor configuration and registration; f. temporally out-of-sequence measurements. Also discussed by Reid, in the section which describes correlating Type-1 sensor datasets with known tracks, is a problem which arises when dealing with non-normally distributed target states.

Kalman Filters and Linear Tracking Models

In general, a Kalman filter implements a Markov process to recursively predict the next state of a linear dynamic system, as a function of the last state, or last set of states. One possible implementation, due to Papoulis (ref. 17), is the following recursive equation:

$$\hat{s}_n = \alpha_n \hat{s}_{n-1} + \beta_n \hat{s}_{n-2} + \gamma_n x_n + \delta_n x_{n-1} \quad (2)$$

This implementation models the next estimated state of the system as a function of the last two estimated states and the most recent two data measurements. Such estimates of system states have proven to be quite useful in domains where a linear model adequately reflects the actual behavior of the system. The estimates are especially appealing for some naval and aircraft tracking applications, in which trajectories manifest little deviation from a linear model. However, when one or more of the sources of nonlinear behavior listed below is manifested, the Kalman filter's predictions may be poor estimates of target location.

In many cases, a target does adhere to a linear track (e.g., a geodesic on the surface of the earth), but there are non-trivial phenomena/stimuli which may cause it to behave nonlinearly. A target's trajectory may behave in a nonlinear fashion when the target is constrained to adhere to a nonlinear path; is attempting to avoid an obstacle; is attempting to pursue a quarry which is moving in a nonlinear fashion; is intentionally trying to deceive an adversary who believes the target to be on a linear path; is acted upon by an external agent who causes the target to deviate from a linear path; or is suffering from an organic malfunction. In such cases, the predictive locational estimates produced by a linear recursive Kalman filter will be off the mark. It can be argued to a certain extent that even though a target trajectory is nonlinear, it appears to be piecemeal linear if a sufficiently small neighborhood is chosen about the sample points. This argument ultimately fails, however, because in practice a sufficiently small neighborhood is not available to the sampling devices, nor is there adequate time for a Kalman filter to adapt to radical deviations from linearity.

Reid's algorithm does not make use of collateral information such as that which might be extracted from map backgrounds with computational geometry software, nor does it avail itself of knowledge concerning the doctrine or organizational structure of opposing forces. Proximity and point-in-polygon queries are now well-solved problems, each with $O[\log n]$ time complexity. There is also no provision for augmenting sensor data with other sources of knowledge, such as the information found in map backgrounds and also in products supplied by the Defense Mapping

Agency such as vectored slope, obstacle, transportation network, hydrology / drainage network, and vegetation overlays.

JUSTIFICATION-BASED TRUTH MAINTENANCE SYSTEMS

The Emergence of Truth Maintenance Technology

In the same year in which Reid's paper appeared in the control theory literature, a paper by J. Doyle appeared in the artificial intelligence literature (ref. 11). With the introduction of justification-based truth maintenance systems, Doyle was making a direct assault on the frame problem formalized by Hayes six years earlier. His solution was a new concept called retraction, in which aged beliefs are marked as false and removed from the valid database. In this way, the number of assertions believed by a reasoning system is constantly changing. In general, the decrementing of evidence for one event is precursory or concurrent with the incrementing of belief for a new event. Doyle's solution to the frame problem was to permit a computer process to retract an instantiated indicator for an event frame back to its original default state, generally presumed to be false (lack of truth is assumed unless there is evidence to the contrary). Therefore, it becomes theoretically possible to revise belief in an event by simply altering one of the logical indicators which provide support for the event. There may or not be a corresponding change of logical state in the set of indicators for some other event in the search space, but the crux of the matter is that belief in an event is either decremented or incremented if new evidence is adopted.

Doyle's concept has come to be known as a *justification-based truth maintenance system (JTMS)*. A *justification* is a logical argument used to derive a new result from a given set of logical assertions. An inconsistency occurs when the negation of one of the given assertions is derived in such fashion, which results in *reductio ad absurdum*. Generally, when an inconsistency is discovered by a JTMS, backtracking is performed to retract an aged assumption, so that only the most current status of an object attribute is kept active by the truth maintenance system. The capability to revise belief in an event by retracting some of the indicators which presage the event permits a truth maintenance system to perform non-monotonic inference, so that the number of facts believed changes as a function of time (Figure 2). According to Doyle, a truth maintenance system serves two roles: as a recorder/consumer of justifications derived by the problem-solving component of a reasoning system, and producer of beliefs in events for which evidence is dynamically changing.




Figure 2. Justification-based truth maintenance system treatment of belief.

All truth maintenance systems to some degree are designed to detect and to record logical inconsistencies in a database. In a temporally dynamic milieu (such as an evolving battlefield environment), an inconsistency indicates that a specific object (force structure element) is no longer in the same location or exhibiting the same behavior demonstrated at some earlier time. At such time, the old reference to the object must be subsumed by the new reference, and is retracted. Note that retraction in general implies three actions in sequence: the posting of new evidence; the detection of inconsistency when the new evidence conflicts with a datum resident to an established database; the resolution of the inconsistency by discarding aged database information in favor of the new evidence.

Event Modeling and Truth Maintenance

If one attempts to discretely model an event such as a flanking maneuver, one is concerned with selecting indicators for the event which point to a flanking maneuver and nothing else. In general, this is very difficult to do, because indicators for events tend to be ambiguous across several events. Nevertheless, the example of Figure 3 enumerates six possible statements to attempt to prove if one is hypothesizing that a flanking maneuver is building over time. The purpose of the three boxes in each one of the indicators is as follows: a white box indicates that no sensor evidence is available to support that indicator; a black box means that evidence is currently available to support that particular indicator; a gray box means that evidence was recently available to support the indicator, but has subsequently been retracted in favor of newer evidence which suggests otherwise. Therefore, at any given time, an indicator for an event may be toggled to one of three discrete positions.




Figure 3. Event frame for a flanking maneuver.

As an example of Doyle's concept in action, consider the hypothetical situation depicted at Figure 4, in which a frontal assault appears to be evolving over time to a flanking maneuver. Time is represented in the vertical direction, with the future towards the bottom of the graphic. Events have been discretely modeled in a Markov fashion, with a frontal assault being modeled to transition to itself, a deliberate withdrawal, a flanking maneuver, or a hasty retreat. In this case, it is apparent that a frontal assault was supported by evidence in the not-too-distant past, as illustrated by the three instantiated indicators and the three retracted indicators. Instantiation is a term from formal logic which translates to "a generalized variable supported by a specific instance of data". At some point in time, the truth maintenance system has received a consensus of evidence from the problem solver which points to a frontal assault as the most likely event. Because of the frame problem, it is not known which of the other three events will occur next, nor is it known if the frontal assault will continue to occur. Note that three of the indicators for a frontal assault have been retracted, although three remain triggered by default. Over time, evidence has correspondingly begun to accrue to support both a flanking maneuver and a deliberate withdrawal. In the example, a higher percentage of indicators are instantiated for a flanking maneuver versus any other event, so if one were forced to make a decision, one could argue strongly (albeit somewhat subjectively) for the flanking maneuver. Several years ago, the author developed a technique called *suspicion accumulation* to address the frame problem (ref. 2). Rather than merely using the percentage of instantiated indicators as a basis of decision, the technique monitors the behavior of the status of indicators as a temporal sequence, as well as the temporal density of evidence for a particular event, as a means of discriminating across all events.




Figure 4. JTMS single best hypothesis selection.

ASSUMPTION-BASED TRUTH MAINTENANCE SYSTEMS

A Reactionary Technology: Assumption-based Truth Maintenance Systems

Shortly after the publication of Doyle's work, concerns arose among logicians and computer scientists alike about the computational time complexity of justification-based truth maintenance systems. Memory requirements with a JTMS are minimal, since only the most credible context is saved at any given moment in time, but the computation required to maintain the best context may be excessively prohibitive, using this variation of depth-first search. Dependency-directed backtracking, of which chronological backtracking is a special form, is a bottleneck. Retraction is a powerful construct, but a practical implementation often suffers from the same prohibitive thrashing behavior exhibited by theorem-proving algorithms in the predicate calculus. Context switching is similarly plagued, since as a high level process, it requires a composition of both chronological backtracking and retraction operations.

Such was the state of affairs with truth maintenance technology when J. de Kleer was a researcher at MIT in the early eighties. At the time, de Kleer was involved in utilizing artificial intelligence to assist in diagnosing and troubleshooting malfunctioning electronic equipment. He experienced frustration when attempting to implement justification-based truth maintenance systems, primarily for the reasons cited in the preceding paragraph. As a remedial action, he decided that it would be useful to permit a JTMS to maintain multiple contexts, while at the same time utilizing a labeling process, to reduce the amount of work required when performing repetitive chronological backtracking. Just as Doyle's paper was a reaction to the challenge issued by Hayes formalization of the frame problem, de Kleer's paper was a reaction to Doyle's approach to truth maintenance. After consulting the literature of his contemporaries (ref. 11, 15-16), de Kleer's solution was assumption-based truth maintenance systems (ATMS), which he claimed were better-suited to deal with inconsistent information, due to their inherent capability to track conflicting hypotheses in parallel (ref. 5-7, 19). He also claimed that his technique avoided retraction and made context switching unnecessary, because context-checking becomes a simple subset test. He later amended these statements somewhat when he realized what he was proposing was actually a form of breadth-first search, and admitted that some sort of control over the ATMS was required (ref. 4). In the graphic at Figure 5, an ATMS scenario starts out with the third context believed more than any other, and at the end the first context has an edge over the other two.




Figure 5. Multiple hypothesis management using ATMS logic.

Partitioning a Reasoning System: Problem Solving vs. Truth Maintenance

De Kleer views a reasoning system as consisting of two components with a duplex feedback loop. This architecture is represented at Figure 6. First, there is a problem solver which performs the usual logical inferences expected of a reasoning agent. The problem solver proves and validates logical assertions and sends the justification (the entire argument, including antecedents, consequents, and implications) to the truth maintenance side. The truth maintenance system in turn is responsible for revising its set of beliefs based on the newly arriving set of justifications, and for sending the results back in efficiently labeled form to the problem solver. The purpose of the lambda-notation in the label is to ensure that the same work is not mindlessly requested of the problem solver.




Figure 6. A reasoning system (de Kleer, 1986).

An ATMS has a built-in capability to maintain multiple hypotheses in parallel, although the search space may become combinatorially prohibitive if the context tree is permitted to branch excessively. Thus control is an issue. Also, an ATMS differentiates between premises, which are the enduring characteristics of a domain, and assumptions, which are tentative beliefs expected to change (ref. 9). When a logical inconsistency is detected during data fusion processing while utilizing a JTMS, it is possible to backtrack in time to an incorrect (or perhaps merely outdated) assumption to retract it, before forward chaining back to the present with a new and consistent line of reasoning. However, this strategy is fundamentally anathema to de Kleer, and he uses it only begrudgingly and sparingly to control the ATMS.

De Kleer in his abstract to reference 5 cites the following features of assumption-based truth maintenance as improvements over traditional truth maintenance:

- a) Facilitates working with inconsistent information;
- b) Includes context-switching as a byproduct;
- c) Avoids backtracking and obviates retraction;
- d) Permits the maintenance of multiple hypotheses in parallel.

As originally conceived by de Kleer in seminal articles published in the AI Journal (ref. 5-7), an assumption-based truth maintenance system makes use of several concepts, to include the following: *premises*; *assumptions*; *justifications*; *contexts*; and *multiple contexts*.

A *premise* is a hard fact about a problem domain, which is not expected to change during the time in which truth maintenance problem solving will occur. For example, it is safe to consider the continent of Europe to be due west and contiguous with the continent of Asia - it is highly unlikely that the relatively long term dynamics of plate tectonics or continental drift will alter this fact during our lifetimes. Any assertion characterized by such an enduring quality over time is considered to be a premise.

An *assumption* is a tentative belief pertaining to a problem domain, which is expected to change over time. Often an assumption consists merely of an assertion which declaratively states the location of some object on a map background. Of course, if the object subsequently moves, the assumption is no longer true.

A *justification* is a logically consistent argument which infers an assertion from a set of premises and assumptions. A justification is concerned with local inference rather than global.

A *context* is a global line of reasoning supported by a logically consistent chain of justifications.

Multiple contexts are sets of plausible but possibly conflicting lines of reasoning.

ATMS Flowchart Symbolology

In lecture notes for a AAAI tutorial on truth maintenance systems, de Kleer presented a convention for flowchart symbols to be used in an assumption-based truth maintenance system. A set of symbols for various ATMS components has been developed, although a complete set is currently lacking. A premise, which in a sense is self-defining since it represents a hard fact about a domain which is unlikely to change over time, is represented as an oval with an arrowhead both sitting upon and pointing at the oval. An assumption is represented as a rectangle. A justification is an oval which always has a set of implication lines leading into it; and may or may not contain implication lines leading out. The logic flow is generally represented on the page from top to bottom, with the symbols at the top involved in implications to derive symbols at the bottom. In practice, entire logical arguments consisting of premises, assumptions, and justifications may be replaced by a label, to concisely depict the global logic for very complex contextual arguments. The labeling process assures that an identifier is installed upon a specific sets of premises and assumptions, with an ordering property, to assure that the truth maintenance system does not repeatedly request the same logical query of the problem solver side of the reasoning system; i.e., the reasoning system "remembers" that it has already performed a specific logical derivation and archives the result.

At figure 7, a premise and an assumption are conjuncted to create a new node, the effect of which is to make a new, simple assertion about troop location in the context of a specific geographic domain. Note that the premise is an enduring characteristic of the environment unlikely to change in the short term, whereas the assumption about troop location is quite likely, perhaps even expected, to change in the short term. The deductive argument, which is actually a case of *modus ponens*, is the justification for creating the new node. In this hypothetical example, the problem solver component of the reasoning system decides that the position of the 3rd engineer company is near enough Delta Run Ford as to be considered the same location, so that the "South of Hill 67" property of Delta Run Ford may be transferred to the Engineer company. It is imperative for the problem solver to somehow deduce the fact that a specific force structure is near enough to a specific stream crossing to be considered "at" it; later it will be shown that computational geometry is useful for this sort of logical validation.




Figure 7. Symbolic conventions for some ATMS components.

Maintaining Parallel, Conflicting Contexts with an ATMS

The power of an assumption-based truth maintenance system derives chiefly from its capability to maintain logical contexts in parallel, even should those contexts conflict with one another. Control of excessive hypothesis generation, and where in the context tree one should be expending effort become the central issues in an ATMS. Since an ATMS is a thinly-guised form of breadth-first search, one must devise a strategy for controlling the "bushiness" of the context tree, or the search combinatorics will be exponential. Certainly, it is not desirable to spawn two new contexts every time a classical statistical test is passed or failed.

Figure 8 illustrates the result of either accepting or rejecting a null hypothesis during a simple statistical test. On the left side of the figure, the null hypothesis, denoted S , is accepted and combined with two premises to derive justification node A . S is meant to represent the adoption of a null hypothesis that two tracks are different from each other, based on statistically different feature vectors. As a result, the truth maintenance system might consider the reported object to be another known emitter, clutter, or an unknown emitter for which no statistics have been gathered. The effect of accepting S is subsequently propagated downwards, to arrive ultimately at node E . At the right side of the figure, the alternate hypothesis, $-S$, is allowed to be logically conjuncted with the same original pair of premises to eventually derive $-E$, which is the negation of the left hand context. $-S$ is the negation of the null hypothesis, which means in this case that the tracks are the same. Both S and $-S$ are expanded in parallel by the truth maintenance system, and combined with other assumptions and premises to arrive at derivations E and $-E$ respectively. If at any time during this process sensor evidence had arrived to support for example $-C$, then the simplest explanation is the right-hand branch, which is the more logically consistent context to maintain with the given constraints.




Figure 8. Two conflicting contexts being maintained in parallel.

UNCERTAINTY IN DATA FUSION PROCESSING

One of the perceived shortcomings of truth maintenance systems as originally conceived is a reliance on strictly logical inputs. For a variety of reasons, logical inputs are notoriously difficult to come by in the domain of data fusion. This is because there are many sources of uncertainty encountered during data fusion processing. A detailed description may be found at reference 20. One source of uncertainty is the resolution error of the sensors used for perceptual-level inputs on the front end of the system. Another is the problem of timeliness and incompleteness: passive sensors in particular may collect information only when targets make themselves known. The algorithms utilized by data fusion systems may make oversimplified assumptions about the probability density functions of the targets, which raises issues about the adequacy of the model. Uncertainty is introduced when we attempt to model a continuous phenomena with a set of discrete indicators, because the frame problem guarantees there will always be instances when a discrete event model does not adequately reflect reality. Finally, algorithms themselves may be probabilistic, in the sense that they may not always be capable of answering a yes-no question with certainty; this latter problem is frequently caused by finite precision arithmetic.

At the perceptual level of intelligence collection, data fusion systems are dependent upon a front end suite of sensors, which function essentially as measuring devices to detect objects in an environment. Any measuring device is subject to error in proportion to the resolution of the device. Different sensors measure different parameters, but there is always an error associated with quantifying information by interpolating between adjacent index points of a measuring tool. Under laboratory conditions, measuring devices invariably perform better than in the field. In practice, it is not uncommon for a system to exhibit an error envelope pattern which is larger than a geographical area of interest. In general, one cannot be assured with one hundred percent probability that a target even resides within the bounds of the envelope pattern.

In addition to sensor resolution error, digitized maps also contain error. First of all, a map is merely a scaled representation of reality. What features are selected for representation is a decision made by the mapmaker, using the results of the most recent set of surveys of a geographical area. Actual features of the area may not be reflected by the map for a variety of reasons, including: surveyor oversight; the fact that by convention the science of surveying does not agree upon including some features; the fact that a survey may be obsolete. Also, in the map digitization process, some features may be washed out completely depending on the scale of the map; whereas others may be represented only semantically but not to scale.

The Need for Robust Hypothesis Testing

According to Huber (ref. 13), any statistical procedure:

1. Should have a reasonably good (optimal or near-optimal) efficiency at the assumed model.
2. Should be robust in the sense that small deviations from the model assumptions should impair the performance only slightly.
3. Should not suffer catastrophically from somewhat larger deviations.

As a case in point, let us consider classical statistical hypothesis testing. We assume a population which behaves with a normal distribution, with a mean of μ and a standard deviation of σ . We will concern ourselves with a two-tailed test, with $\alpha = .025$. The Z-values corresponding to an α of .025 are $z = -1.96$ and $z = 1.96$. If our sample statistic is less than -1.96 or greater than 1.96 , then the statistical test of means is failed, which means that we should reject the null hypothesis that the sample is different from the one for which statistics have been gathered. The critical region for the test is depicted at Figure 9.

$$n(x; \mu, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{1}{2}\left[\frac{(x-\mu)}{\sigma}\right]^2}, \quad -\infty < x < \infty \quad (3)$$




Figure 9. Classical statistical hypothesis testing.

Classical statistical testing is brittle; i.e., it is non-robust. A change in a single sample value can cause a passed test to be failed, or a failed test to be passed. This phenomenon may be described as the "near-miss, barely-hit" paradox (figure 10). The cause of the brittleness of the test is the insistence upon a rigid confidence region, such as five percent or ten percent of the total distribution. These arbitrary percentages are chosen independently of the observed population, which for various reasons (such as maneuvering) may contain a larger variance from the mean than that computed with the traditional formula, which relies upon a static population.




Figure 10. Statistical brittleness: the near-miss, barely-hit paradox.

A proposed solution to the traditional brittle test is a robust one designed in the following fashion. Rather than insisting upon a Z-statistic which corresponds to a specific confidence level, we relax the condition and assume that a Z-statistic itself may be normally distributed about its true value (figure 11). This gives rise to two new distributions, with means respectively equal to the negative or positive Z-statistics at the critical juncture, and variances chosen as a function of the thickness and breadth of the tail (kurtosis) of the parent distribution.




Figure 11. Permitting critical values themselves to be normally distributed.

The purpose of the proposed robust statistical test is to selectively control the generation of hypotheses by an assumption-based truth maintenance system. One could for example, push a new context onto the stack if a statistical test is passed or failed based upon a sample statistic falling within one standard deviation of the mean critical value selected for the test. Otherwise, the new context is not created. Clearly, it is preferable to postulate a new hypothesis only when a statistical test is most ambiguous; i.e., when the sample statistics are within a small neighborhood of the critical value, which subscribes to Huber's definition of robustness.

The robust strategy to spawn a new context only when a statistical test is barely passed or barely failed could do much to control the sheer size of a context tree when maintaining multiple hypotheses with an assumption-based truth maintenance system. Of course, the strategy could also be used in Reid's tracker to control search for his multiple-scan algorithm, which is the analogous multiple hypothesis technique in the discipline of control theory. Even more search constraints are available to be exploited if one turns to the relatively recent field of computational geometry, which leads us to the next section.

COMPUTATIONAL GEOMETRY AND LOGICAL INPUTS FOR TRUTH MAINTENANCE

Advances in computational geometry (CG) research made since the Reid paper appeared may now make it feasible to treat some of the issues which were not feasibly addressable in 1979. In particular, some of the nonlinear tracking problems may be approachable with the new technology. Optimized algorithms for proximity and point-in-polygon testing may supplement a Kalman filter with a constraint set to produce a nonlinear estimate of target location. These low-level functions may serve to answer some of the perceptual level queries which can for example, probabilistically place a target on a roadbed, vs. not having any collateral knowledge about the target. If one knows that a target is a wheeled vehicle constrained to remain on a roadbed, then this knowledge can be used to constrain estimates of target location to lie on a nonlinear path, if in fact the road trajectory happens to be nonlinear.

Figure 12 illustrates two of the computational geometry algorithms developed by the Signals Warfare Directorate to solve the proximity and point-in-polygon problems for tactical maps (ref. 3). The equidistance locus is a vector-based approach to the map reasoning problem, and relies upon the actual distribution of separates on the map to solve the proximity problem. The inclusion issue is solved by dropping the normal vector from a query point to a boundary, and observing to which side of the boundary the vector points. These techniques can within the resolution of a sensor, advise a reasoning system which of a network of roads is nearest to a sensor-derived query point, or whether the point is inside an area of interest to a tactical commander. Therefore, the geometric results might be used as primary inputs to a truth maintenance system seeking to instantiate slots in an event frame, or as collateral inputs to Reid's algorithm to check if the Kalman estimates are reasonable, given the map at hand.

The nearness and inclusion algorithms are currently being adapted to a 1:50,000 scale map of Killeen, Texas, for which over ten megabytes of vectorized interim terrain data (ITD) is available, as well as an eight megabyte CD-ROM formatted raster map background. The hardware suite to which the algorithms have been ported is a Macintosh IIx, which will have thirty two megabytes of RAM when fully loaded. The map, which is a digitized version of the hardcopy printed by the Defense Mapping Agency, will be displayed on an E-machines T19 color monitor. In addition, the author has procured from Rockware, Inc. a software tool called MacGridzo, which is a geographic information system tool designed to produce topographical contours from a list of control point elevations provided by data available from the United States Geological Survey

(USGS) organization. The research plan is to simulate movement across the map terrain, with some movement constrained to the transportation network. A Kalman filter will be used to predict the new locations of targets on tracks, while at the same time the computational geometry algorithms will be used to compute how well predicted tracks are adhering to the map transportation network.




Figure 12. Proximity and Inclusion Algorithms

Justification-based Truth Maintenance Systems and Zero-scan Algorithms

Reid introduces a zero-scan algorithm to produce the single best hypothesis after processing each data set. The concept of selecting the single best hypothesis is superficially analogous to that espoused by Doyle's justification based truth maintenance systems. The analogy is only approximate because of the kinds of reasoning utilized by the systems. Reid's algorithm uses the principle of maximum likelihood to select the best hypothesis, whereas Doyle seeks logical inconsistency in a database, which is implemented with a chronological backtracking and retraction capability. When Reid's system does not fuse together measurements which should be fused, he suggests opening up the covariance matrix as a remedial action. Doyle's system in effect throws away the averaging leverage of old statistics by retracting the fact that an object was seen in the past at position p, and is now in light of new evidence believed to be at position q. The object will, according to Doyle, continue to remain at q until evidence to the contrary is forthcoming. Such is the behavior of a belief system based on defeasible logic. Reid uses Kalman filter estimates to predict future positions of a target. Doyle's system proceeds from logical assertions about the locations of objects, and therefore has no statistical estimates for new locations of an object. However, Doyle knows where the object could logically go, based on collateral evidence concerning the trafficability of semantic features in the vicinity of the last known position of the

object, and upon a description of the locomotive capabilities of the object. Reid does not avail himself of such information. It can be seen that the Reid and Doyle philosophies to problem solving share certain concepts, yet each contributes certain advantages that the other does not.

Assumption-based Truth Maintenance Systems and Multiple-scan Algorithms

In Reid's paper, the analog of the multiple hypothesis capability of an assumption-based truth maintenance system is the multiple-scan algorithm. The multiple scan algorithm permits a user to build several competing statistical hypotheses in parallel, with the stipulation that each hypothesis is produced by comparing the predictive locations of a Kalman filter with actual target measurements. The same concepts as discussed in the comparison of zero-scan algorithms to JTMS apply here, except that now parallelism is permitted.

UNIFYING THE CONTROL THEORY APPROACH TO MULTIPLE TRACKING WITH A COMPUTATIONAL GEOMETRY-DRIVEN TRUTH MAINTENANCE APPROACH

Table 1 is a side-by-side comparison of the attributes of Reid's approach from classical control theory with that of Doyle's and de Kleer's from the world of artificial intelligence and logic. Reid deals with uncertainty by appealing to the theory of probability, while the world of truth maintenance utilizes real-time computational geometry inputs to compute solutions to logic-based proximity and inclusion queries. In control theory, statistical clustering has been traditionally used to recognize events via templating, whereas in truth maintenance systems, Boolean slot activation is used to trigger belief in event models.

Table 1. Multiple target tracking: Control theory vs. AI / logic

Feature/System	Reid's System	Doyle's TMS	De Kleer's ATMS
Style of Reasoning	Probabilistic	Enhanced Logic	Enhanced Logic
Predicted Track	Kalman filter	Logical constraints	Logical constraints
Track Association	Maximum likelihood	Comp. Geometry	Comp. Geometry
Best Hypothesis	Zero-scan Algorithm	Dependency-directed back	Most consistent context
Multiple Hypothesis	Multiple-scan Algorithm	N/A	Conflicting contexts
Collateral inputs	Does not use map	Uses map constraints	Uses map constraints
Event recognition	Statistical clustering	Boolean slot activation	Boolean slot activation

Figures 13a-b contain examples to demonstrate the utility of enlisting the aid of a map when predicting the new positions of targets. Figure 13a shows hypothetical target location estimates as predicted by a recursive Kalman filter; whereas figure 13b shows possible corresponding estimates made by computational geometry algorithms. Refer to figure 13a for a moment. In track 1, the target is predicted to climb a fairly steep hill in as little time as it took to traverse level ground. In track 2, the Kalman filter looks linearly ahead, and ignores logical constraints suggesting that the target is navigating a road. In track three, the target, which had previously been tracked over land, suddenly plunges into a lake. Compare to the CG-predicted locations shown in figure 13b. What is meant to be suggested by this graphic sequence is that Kalman filter estimates might be enhanced by the logical inputs provided by a computational geometry component of a multiple target tracker. This is not intended to demean the utility of a Kalman filter, but simply to suggest that the two techniques acting in collusion might provide a more powerful tracking capability to a tactical commander than either technique acting by itself.




Figure 13a. Kalman filter estimates of target location for three tracks.




Figure 13b. Computational geometry estimates of target location for the same tracks.

Returning to de Kleer's reasoning system concept, it can be seen that we might productively borrow from both control theory and computational geometry to produce a multiple target tracking capability in a reasoning system. The problem solver side of the system may tap both Reid's technique and CG-derived collateral map inputs to derive justifications which are sent to the truth maintenance side. In turn the truth maintenance system performs the operating and housekeeping chores of turning event slots on and off, retracting aged events, and maintaining multiple, conflicting hypotheses in parallel.




Figure 14. A multiple target tracking reasoning system.

SUMMARY

We have proposed a multiple target tracking system which combines positive attributes of three separate disciplines: control theory, AI / logic, and computational geometry to arrive at a composite tracker more comprehensive than that which any single discipline acting alone can provide. One of the features of the system is that in the absence of collateral map inputs, the system may run in a streamlined mode and utilize Reid's tracker, which does not require map inputs for its estimation equations. However, when we may avail ourselves of the map, computational geometry imposes a system of logical checks and balances upon the statistical tracker to improve upon the estimates achieved by Kalman filter logic alone.

The technology of assumption-based truth maintenance systems (ATMS) appears to be viable for arbitrating among multiple contexts during data fusion processing. An effort has been made to link the various components of an ATMS to analogous components of a probabilistic multiple target tracker developed by D.B. Reid. When hampered by a lack of collateral information, the ATMS works in a streamlined mode and essentially reduces to Reid's system. However, when non-sensor data is exploitable, the ATMS may be capable of improving upon the Kalman filter location estimates predicted by the statistical system in standalone mode. In some

cases, the collateral information maintained by the ATMS may serve to avoid the type I and II statistical errors which thwart tracking systems which do not avail themselves of the extra information. An additional cost-reduction benefit of the ATMS is the way in which logical justifications are recorded with specific labels during data fusion processing. The unique labeling process ensures that work is not needlessly repeated to rederive old results. Some of the factors circumvented by Reid in his paper are shown to be addressable by the technology of truth maintenance systems. At the conclusion of the paper, the issue of uncertainty in data fusion processing is treated, and it is suggested that the pure logical inputs anticipated by truth maintenance system technology might be enhanced with an uncertainty calculus. Specifically, it is demonstrated that Huber's robust statistics is a suitable technology to permit relaxation at the juncture between the critical and acceptable regions in classical statistical hypothesis testing. This relaxation process can control excessive generation of multiple hypotheses for either an assumption-based truth maintenance system, or for that matter, Reid's multiple-scan algorithm. The process also permits the generation of hypotheses which might not have been produced with a conventional statistical system.

ACKNOWLEDGMENTS

This paper is in large part a response to a perceptive request made by Frank White, chairperson of the Joint Directors of Laboratories, Data Fusion subpanel, to compare the multiple hypotheses capability of assumption-based truth maintenance systems vs. a technique espoused by D.B. Reid in the control theory literature; I thank Mr. White for the suggestion. I am indebted to a large number of scientists and engineers working in a variety of disciplines for insights which arose and were pursued during the course of this research. In particular, I would like to thank Patrick Hayes and Henry Kyburg for a productive research contract with my organization which related the frame problem and non-monotonic reasoning to the Army's concept of situation assessment. Also instrumental to the progress of the work has been an excellent working relationship with Dave Hislop of the Army Research Office, who has brought to my attention several excellent efforts in uncertainty management which he has personally directed. I would also like to cite Rudolf Kalman for an insightful talk entitled "Identification of Systems from Noisy Data - A New Look at Statistics from the Real World", which he presented at the Ninth Army Conference on Applied Mathematics and Computing, held at the University of Minnesota in June of 1991. Finally, I would like to thank Ray Freeman of the ASAS program office, Robert Somoano of the Jet Propulsion Laboratory, James Allen and Alan Gernalnick of the Engineering Topographic Laboratory, and Dave Grubb, Chris Bogart, Doug Chubb, and Richard Antony of my organization for being technically receptive to the vectorized map reasoning tools developed in-house with computational geometry research. It is hoped that these tools may ultimately provide logical inputs from the problem solver side of a reasoning system to the truth maintenance side, which seeks to validate or refute logic-based indicators that point to different contexts.

BIBLIOGRAPHY

1. R. Antony, "Eight Canonical Forms of Fusion", CECOM Center for Signals Warfare report (1991).
2. T. Cronin, "The Combinatorics of Message Filtering", Proceedings of the Thirty-Second Conference on the Design of Experiments in Army Research, Development, and Testing, Army Research Office Report 87-2, Monterey CA (1986).

3. T. Cronin, "Dynamic Coordinate-to-Feature Association for Command and Control Applications", Technical Proceedings of the 1988 Tri-service Data Fusion Symposium, Vol. 1, Naval Air Development Center, Warminster PA (1988).
4. J. de Kleer "A Comparison of ATMS and CSP Techniques", Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, V. 1, Detroit MI (1989), pp. 290-296.
5. J. de Kleer, "An Assumption-based Truth Maintenance System", Artificial Intelligence 28 (1986), pp. 127-162.
6. J. de Kleer, "Extending the ATMS", Artificial Intelligence 28 (1986), pp. 163-196.
7. J. de Kleer, "Problem Solving with the ATMS", Artificial Intelligence 28 (1986), pp. 197-224.
8. J. de Kleer, "Exploiting Locality in a TMS", Proceedings of the Eighth National Conference on Artificial Intelligence, Boston MA (1990), pp. 264-271.
9. J. de Kleer and K. Forbus, "Truth Maintenance Systems (or Consistency) Management", AAAI-86 Tutorial no. 21 Lecture Notes, University of Pennsylvania (1986).
10. J. de Kleer and B.C. Williams, "Back to Backtracking: Controlling the ATMS", Proceedings of the Fifth National Conference on Artificial Intelligence, Philadelphia PA (1986), pp. 910-917.
11. J. Doyle, "A Truth Maintenance System", Artificial Intelligence 12 (1979), pp. 231-272.
12. P.J. Hayes, The Frame Problem and Related Problems in Artificial Intelligence, in: "Artificial and Human Thinking", A. Elithorn and D. Jones, eds., Josey-Bass, San Francisco CA (1973).
13. P. Huber, "Robust Statistics", John Wiley and Sons, New York NY (1981).
14. R.E. Kalman, "A New Approach to Linear Filtering and Prediction Problems", J. Basic Eng., Transactions of ASME Vol 82-D (1960).
15. D. McAllester, "Reasoning Utility Package User's Manual", MIT AI Laboratory Report AIM-667, Cambridge MA (1982).
16. D. McDermott, "Contexts and Data Dependencies: a Synthesis", IEEE Trans. Pattern Analysis and Machine Intelligence 5 (3) (1983).
17. A. Papoulis, "Probability, Random Variables, and Stochastic Processes", McGraw-Hill Book Company, New York NY (1965).
18. D.B. Reid, "An Algorithm for Tracking Multiple Targets", IEEE Trans. Automat. Contr., Vol. AC-24, No. 6 (1979), pp. 843-854.

19. R. Reiter and J. de Kleer, "Foundations of Assumption-based Truth Maintenance Systems: Preliminary Report", Proceedings of the Sixth National Conference on Artificial Intelligence, Seattle WA (1987), pp. 183-188.

20. E. Waltz and J. Llinas, "Multisensor Data Fusion", Artech House, Inc., Boston MA (1990).