

EVIDENTIAL REASONING APPLIED TO SINGLE-OBJECT LOSS-OF-CUSTODY SCENARIOS FOR TELESCOPE TASKING

Andris D. Jaunzemis*, Marcus J. Holzinger†

Evidential reasoning and modern data fusion models are applied to the single-object loss-of-custody scenario in ground-based tracking. Upon a missed observation, the cause of non-detection must be quickly understood to improve follow-up decision-making. Space domain awareness (SDA) sensors, including a brightness sensor and an All-Sky camera with an optical-flow-based cloud detection algorithm, are conditioned as Dempster-Shafer experts and used to assess the cause of a non-detection. Telescope re-tasking is also approached using Dempster-Shafer theory by planning the next observation to minimize an estimated lack-of-information. Results from real-world operational sensors show the algorithm's ability to adjust to changing observation conditions and re-task the primary electro-optical sensor accordingly.

INTRODUCTION

Space situational awareness (SSA) is concerned with accurately representing the state knowledge of objects in the space environment to provide better prediction capabilities for threats such as potential conjunction events. More recently, the discourse on SSA has turned toward space domain awareness (SDA), reflecting the ever-growing reality of world-wide space capabilities and the impact that decisions in the space environment can have on a global relational scale. The space community as a whole suffers from a problem of producing high quantities data (in the form of tracks) but being unable to produce significant data on any specific object or event to increase understanding of that event. Currently, there are over 20,000 trackable objects in the space object catalog.¹ Due to observational constraints imposed by orbital mechanics, the limited number of space-observing sensors are unable to observe each object. This hinders the ability to reliably provide information on maneuvers or other events in space. Therefore, more emphasis is being placed on algorithms and processes that have an ability to ingest disparate data from many sources and fuse an understanding of the greater situation of the space domain.

In typical Bayesian reasoning, deterministic probabilities are placed on event hypotheses under the assumption that the only possible realizations of this hypothesis are true or false. However, in complex decision-making contexts, information is not always best-represented in this strictly binary manner, since some evidence for a particular hypothesis might also involve ambiguity. An expert might be able to confirm or refute a given set of hypotheses, but it cannot attribute belief to any hypotheses for which it is not an expert. For this reason, evidential reasoning methods, such as Dempster-Shafer theory, quantify this ambiguity in situation knowledge, leading to more realistic modeling of human analyst processes.²

*Graduate Research Assistant, Georgia Institute of Technology, Atlanta, GA, AIAA Student Member

†Assistant Professor, Georgia Institute of Technology, Atlanta, GA, AIAA Senior Member

Autonomous decision making processes are typically complex and multi-faceted, requiring rigorous definition of information flow between different analyses to provide a higher-quality representation of the true state. A particularly common defense-related model, the Joint Directors of Laboratories (JDL) data fusion model, provides a high-level mapping of the different process-level components that are involved in autonomous decision-making.^{3,4} Instead of prescribing a hierarchical data flow, JDL encourages an organization that permits data-sharing between different components. The revised version of JDL, called the Data Fusion Information Group (DFIG) model, also allows user input to be considered.

The loss-of-custody scenario is a candidate SDA application for both data fusion and evidence theory. When tracking a space object using a ground-based telescope, if the track disappears from the field of view, a logical hypothesis is that the spacecraft has maneuvered since the last attempted observation. If this is the case, timely follow-up detection is critical in detecting and reconstructing the maneuver, requiring re-tasking of sensors. However, it is also possible that the observation conditions have temporarily degraded, perhaps due to local weather conditions or proximity to bright objects such as the moon. In this case, alerting and re-tasking sensors to look for the missing object causes a loss of information from other targets. This places a premium on correctly identifying whether a maneuver is the cause of the loss-of-custody. A better understanding of the whole observation environment, fusing information from multiple SDA processes, can provide insight on likelihood of each scenario and allow for better modeling of the decision process.

This work begins by introducing concepts of evidential reasoning and data fusion to describe how they can be used to better model a realistic decision-making process. Then, contributions to SDA are provided through 1) the application of the JDL/DFIG paradigm to autonomous telescope tasking for the loss-of-custody scenario, 2) the application of rigorous Dempster-Shafer evidential reasoning to loss-of-custody processes, 3) the conditioning of available SDA sensors into Dempster-Shafer experts, 4) a proposed methodology for telescope re-tasking based on reducing the gap between belief and plausibility, and 5) testing of this loss-of-custody scenario framework using operational data from hardware at the Georgia Tech observatory.

BACKGROUND

This section lays the groundwork for the novel theory developed in the following sections, beginning by introducing Dempster-Shafer theory, a well-known formulation of evidential reasoning. This is followed by discussion of data fusion techniques, focusing on JDL/DFIG.

Dempster-Shafer Basics

In typical probabilistic approaches, a precise probability is assigned to an event, regardless of the quality of the data.⁵ This approach is sensible when considering the truth value of a hypothesis: the hypothesized event either occurred (is true), or did not occur (is false). Bayesian theory commonly represents these probabilities using the pair (p, q) respectively, with $p + q = 1$. In contrast, theories of evidence and other so-called possibilistic approaches add an extra dimension to the expression of uncertainty, using a probability triple (p, q, r) to model the categories “known to be true,” “known to be false,” and “don’t know” for each proposition.² These imprecise probabilities introduce indeterminism to decision-making analyses, but this loss of precision can be viewed as a strength in that it can more faithfully represent the reality of the decision environment.⁵ While it remains true that each statement is either true or false, the analyst can only form a decision based on available evidence supporting or refuting articulated hypotheses. The analyst can leverage available evidence

to produce levels of certainty for (p) and against (q) each proposition, leaving $r = 1 - p - q$ to quantify residual ambiguity.² Evidence theory improves decision-making in high-consequence systems through its ability to better characterize subjective belief.⁶

Dempster-Shafer theory, also known as the mathematical theory of belief functions, is a well-known evidential theory framework with demonstrated usefulness in engineering applications for the representation of epistemic uncertainty and risk analyses.⁶ Dempster-Shafer theory formalizes the use of available evidence to attribute belief to sets of hypotheses for decision making. The following terms and mathematical notation are common in Dempster-Shafer formulations, using typical evidential reasoning notation from Dempster.²

The frame of discernment, Θ , is a mutually exclusive, collectively exhaustive set of discrete hypotheses $\theta \in \Theta$. Mutually exclusive refers to the fact that only one hypotheses, θ_i , may occur at a time: $\mathbb{P}[\theta_i] \cap \mathbb{P}[\theta_j] = 0 \forall i, j \in \Theta, i \neq j$. Collectively exhaustive means that together the frame of discernment encompasses the full spectrum of possibilities so that one of the hypotheses, θ_i , must be true.

$$\Theta = \{\theta_1, \theta_2, \dots\} \quad (1)$$

The power-set of the frame of discernment, 2^Θ , forms all possible disjunctive combinations of the elements in set Θ . Note that, if Θ consists of n elements, the power-set consists of $2^n - 1$ elements; in classical Dempster-Shafer theory, belief mass cannot be assigned to the empty-set (since the frame of discernment is assumed to be exhaustive). Also note that $\Theta \subset 2^\Theta$. The set Θ is also referred to as the truth-set since it represents the disjunctive combination of every element in the mutually exclusive and collectively exhaustive frame of discernment, meaning one of the propositions in this set must have occurred.

A basic belief assignment (BBA) represents an expert's belief in each hypothesis based on the evidence available to that expert. The BBA for the i^{th} expert is given by a belief mass function $m_i : 2^\Theta \rightarrow [0, 1]$.

For ease of discussion and use, a number of useful BBAs are typically defined. In a vacuous BBA, all the belief mass is assigned to the truth-set, Θ , such that $m_i(\Theta) = 1, m_i(A) = 0 \forall A \subseteq 2^\Theta \setminus \Theta$. A simple BBA is one in which the focal set, or the set of hypotheses with non-zero belief mass, consists of only two elements: the truth-set and one other hypothesis, as in $m_i(A) = p, m_i(\Theta) = 1 - p, m_i(B) = 0 \forall B \in 2^\Theta \setminus \{A, \Theta\}$.

The notions of belief and plausibility form lower and upper bounds on the probability that a given proposition is provable from the available evidence. Belief and plausibility can be computed from BBA m_i using Eqs. (2) and (3) respectively.

$$\text{bel}_i(A) = \sum_{B \subseteq A} m_i(B) \quad (2)$$

$$\text{pl}_i(A) = \sum_{B \cap A \neq \emptyset} m_i(B) = 1 - \text{bel}_i(\neg A) \quad (3)$$

where $\neg A$ is the negation or complement of hypothesis A . In other words, expert i 's belief in, or support for, hypothesis A is composed of the sum of the belief masses attributed to A and its subsets, whereas its plausibility of hypothesis A is composed of the sum of the belief masses attributed to any hypothesis whose intersection with hypothesis A is non-empty. Notice that the alternate equation

of plausibility in Eq. (3) gives a relationship to the belief of the complement of a hypothesis. It is worth noting that, since the truth-set Θ represents the disjunctive combination of an exhaustive set of hypotheses, the belief and plausibility of the truth-set must both be equal to 1, a useful fact for checking implementation issues.

Numerous methods exist for combining BBAs from multiple experts to form a fused mass function.^{7,8} Each method exhibits slightly different properties, so implementation should take into consideration use-cases of this fused belief and characteristics of the evidence sources. A common BBA combination technique is Dempster’s conjunctive rule, which is commutative, associative, and admits the vacuous BBA. Dempster’s conjunctive rule of combination, shown in Eq. (4), is often represented using the \oplus operator.

$$m_{i\oplus j}(A) = (m_i \oplus m_j)(A) = \frac{\sum_{B\cap C=A} m_i(B)m_j(C)}{1 - \sum_{B\cap C=\emptyset} m_i(B)m_j(C)} \quad \forall A \subseteq \Theta \quad (4)$$

The term in the denominator handles conflict between the bodies of evidence. Some uses of Dempster’s rule lead to counter-intuitive results in the presence of extreme conflict, an observation typically referred to as Zadeh’s paradox.⁹ However, the scenario in Zadeh’s paradox can be resolved by more carefully adhering to Cromwell’s Rule, i.e. not assigning a probability of exactly 0 or 1 to any particular prior.¹⁰ This caveat, with the inclusion of the open-world assumption, i.e. admitting that the actual true event might lie outside the theorized set of possible events, led to the development of the Transferable Belief Model as a derivative of Dempster-Shafer theory.¹¹ The constraints of this application allow the classical Dempster-Shafer implementation to be appropriate.

Another important note about Dempster’s rule is that it does not possess the property of idempotence. Subsequent evidence is assumed to be statistically independent of previous evidence. Therefore, when using Dempster’s rule the evidence must be assumed to be distinct; otherwise, repeated evidence will be heavily weighted. Alternate combination rules have been developed that do enforce idempotence, which can be employed in the case of non-distinct bodies evidence.¹²

For a more complete discussion on important developments in Dempster-Shafer theory, Yager and Liu compiled a book of classic works, reviewed by Dempster and Shafer, on the theory of belief functions.¹³ With Dempster-Shafer theory in mind, a data fusion framework should be selected to construct an algorithm that utilizes Dempster-Shafer experts.

Data Fusion using JDL/DFIG

The Joint Directors of Laboratories (JDL) data fusion model is a framework for planning and visualizing information flow within a complicated system. Its more-recent revision, the Data Fusion Information Group (DFIG) model, involves six levels of processing, ranging from sub-object level signal processing through course-of-action impact assessment and process refinement for re-tasking assets based on the proposed course of action.³ These levels are detailed in Fig. 1. A system that implements data association and estimation events of all levels will permit better understanding of entities in complex systems.⁴ The ordering

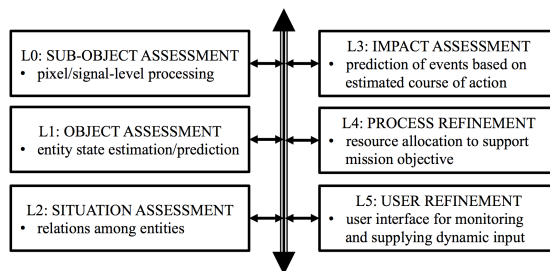


Figure 1. JDL/DFIG process level descriptions⁴

of the levels in the JDL/DFIG model was never intended to imply a hierarchy of processes.⁴ Rather, the model is typically represented with each level sharing a common data bus, as indicated in Fig. 1. For instance, object and sub-object processing and estimation events can benefit from knowledge from other low-level sensors as well as the overall situation assessment to improve performance.

In SDA, many different phenomena are exploited for detection or estimation processes on various pieces of hardware, occupying every level in the JDL/DFIG framework. For instance, taking a noisy image and performing dark-frame subtraction and sub-pixel object identification techniques could be classified as L0, while the classification of these objects as stars or SOs requires L1 capability and associating these objects together for correlation is an L2 task. SDA applications also highlight the non-hierarchical nature of the different process levels. Re-tasking a telescope to look in an area, an L3/L4 task, can prime an L0/L1 object identification algorithm with the commanded pointing to aid SO or star identification. Similarly, Dempster-Shafer experts exemplify the non-hierarchical framework through the combination of evidence from a variety of sources to form a fused situational assessment. In this work, L0/L1/L2 algorithms for SO detection, cloud cover estimation, and sky brightness evaluation are fused with L3/L4 estimates of re-tasking decisions to utilize sensor resources more effectively.

THEORY

In this section, novel theory regarding the contributions is developed. First, decision-making implications of quantified residual ambiguity are discussed to provide intuition for its usefulness in autonomous algorithms. Then the loss-of-custody scenario and algorithm is described. An evidential reasoning hypothesis testing approach to non-detection assessment is developed, after which the re-tasking approach is discussed.

Belief-Plausibility Gap

The belief-plausibility gap, the difference between belief and plausibility, represents a lack of knowledge of the truth given by the available evidence. Since belief mass is only assigned to hypotheses based on direct evidence, the remainder of the belief mass is attributed not to the negation of that hypothesis (as in a Bayesian scheme) but instead to the truth-set. This avoids falsely attributing evidence that the expert really does not have, better representing realistic decision-making processes. This lack of knowledge is what Dempster calls residual ambiguity,² but is also more colloquially referred to as a degree of ignorance. The ignorance associated with knowledge of hypothesis A for expert i is given in Eq. (5).

$$\text{ig}_i(A) = \text{pl}_i(A) - \text{bel}_i(A) \quad (5)$$

Belief and plausibility form the lower- and upper-bounds of a probability interval, bounding the precise probability of a hypothesis. As ignorance in a particular hypothesis approaches zero, the system approaches a state of sufficient evidence to reduce to Bayesian probability calculations for that hypothesis. Indeed, a Bayesian BBA is one in which the focal set elements are all singleton hypotheses.

In order to improve decision-making capabilities, one can focus on taking actions that maximally reduce ignorance in any or all hypotheses. In doing so, the bounds on the precise probabilities will shrink and the true state can be determined with less ambiguity. Reduced ambiguity aids

decision-making by providing a clearer picture of the truth of the situation, ensuring appropriate actions are taken. Therefore, the autonomous algorithm developed in this paper focuses on gathering information-rich, low-ignorance data, aiming to reduce the gap between belief and plausibility.

Loss-of-Custody Algorithm

The scenario of interest, termed the loss-of-custody scenario, involves the use of a ground-based telescope and primary electro-optical sensor to detect a space object. A number of other sensors contribute observation environment information, such as sky cover and background sky brightness.

For this scenario, the telescope is tracking a space object from the space object catalog. On its next attempted observation, the target is no longer found in the telescope’s field of view. The tracking algorithm must now determine the cause and re-task the telescope as necessary. One hypothesis for the loss-of-custody is a maneuver executed by the spacecraft since the previous observation. In order to accurately characterize and understand the maneuver, as well as update the space object catalog, a follow-up observation is desired in a timely manner. Depending on the priority of the target, this could require re-tasking of the telescope as soon as possible or even alerting other resources. However, another possibility states that the spacecraft has not actually maneuvered, but the imaging sensor was unable to obtain enough signal to identify it in the image. This could be due to a number of causes, including local weather conditions or proximity to the moon, that cause poor observation conditions. In this case, sending alerts and re-tasking assets to search for a maneuver constitute a waste of resources and could lead to loss of data. A preferred response might be to temporarily re-task the telescope to another object until observation conditions improve.

For this inaugural research effort, the authors impose a “single-object world” restriction. This means that for the purposes of this study only one space object exists, the target space object, which significantly restricts the decision space by avoiding correlation questions.




Figure 2. Loss-of-Custody Algorithm Block Diagram and JDL/DFIG Relationships

The algorithm flowchart for this task is described in Fig. 2, beginning with a command to slew the telescope to the expected target location for imaging. The image is processed to identify stars and space objects in the image. Note that, for the present study, there is only ever assumed to be at most one space object (a single-object world assumption), to avoid the problem of association for now. Given the single-object world assumption, there are a limited number of possible outcomes from the space object identification process: a) the space object is found in the expected location, b) the space object is found, but not in the expected location, or c) the space object is not found within the image. The first outcome indicates a successful tracking observation, and the new observation can be used to update the state estimates in the catalog. The second outcome directly indicates that an anomaly (e.g. a maneuver) has occurred since the last track update since the object is found in an area not predicted by quiescent propagation. This potential situation will be avoided for this particular study but is an interesting area for future extension, perhaps using a control cost or Mahalanobis distance

metric for maneuver detection.^{14,15} The third outcome is of particular interest to this paper. The space object is not found anywhere in the image, so without direct evidence of an anomaly (as in case 2), a hypothesis test must be performed to determine whether it is likely that an anomaly caused the non-detection. This hypothesis test is approached using a Dempster-Shafer formulation of evidential reasoning, resulting in a plausibility that the non-detection is due to an anomaly.

If the anomaly plausibility is high, the telescope could attempt to re-acquire it by searching the reachable area of the spacecraft. Re-acquisition would provide direct evidence of a maneuver, reducing the belief-plausibility gap. Alternately, if the anomaly plausibility is low, the telescope could continue unchanged to its next planned observation attempt, seeking to verify that a maneuver has not occurred to also reduce the belief-plausibility gap. The re-tasking approach will be discussed at length in a later section.

Each process in the flowchart in Fig. 2 is denoted with a number corresponding to the appropriate level in the JDL/DFIG formulation. The algorithm endeavors to fuse low-level pixel-by-pixel processing with high-level decision-making functions to aid each other. Figure 3 includes high-level details of the processes involved in this algorithm and how they line up with the JDL process level specifications from Fig. 1. Since this is an exploratory study on the applications of JDL/DFIG and evidential reasoning to the telescope tracking problem, the added complexity of a user interface (L5) is not yet considered.

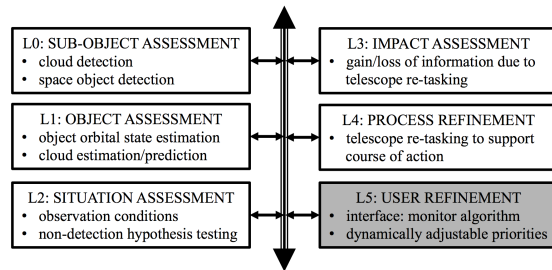


Figure 3. JDL for proposed tracking algorithm

Since this is an exploratory study on the applications of JDL/DFIG and evidential reasoning to the telescope tracking problem, the added complexity of a user interface (L5) is not yet considered.

Non-Detection Hypothesis Testing

To assess whether the target object should have been visible, a hypothesis test must be performed to determine plausible causes of non-detection. Of particular interest is determining whether a spacecraft maneuver is a cause for the non-detection, which may warrant sensor re-tasking to search for the object. Here, Dempster-Shafer evidential reasoning is employed to fuse different sources of evidence and assess the plausibility of an anomaly such as a maneuver.

The first step in implementing Dempster-Shafer reasoning is to define a mutually exclusive and collectively exhaustive set of hypotheses, a frame of discernment, that covers the meaningful causes of non-detection. There are a number of potential causes for a missed detection apart from a spacecraft maneuver: 1) target object obstructed by local weather conditions, 2) foreground sky brighter than target object, 3) target object too near to other bright celestial object, particularly the moon, 4) poor geometry between target object, the Sun, and the observer, diminishing the target’s optical signature (e.g. spacecraft in eclipse or poor phase angles for illumination), or 5) target object coincident with another space object, a planet, or a star. Each of these causes can be thought of as a separate piece of evidence against the hypothesis that an anomaly is the cause of non-detection. Since these all contribute evidence to a common event (i.e. non-anomaly causes of non-detection), they are not represented by separate hypotheses but instead by a straightforward binary hypothesis scheme: non-anomalous causes (N) and anomalous causes (A). Therefore, the frame of discern-

ment can be represented by the following set:

$$\Theta = \{N, A\} \quad (6)$$

The fully enumerated power-set is then represented by the following set:

$$2^\Theta = \{\{N\}, \{A\}, \{N, A\}\} = \{\{N\}, \{A\}, \Theta\} \quad (7)$$

These hypotheses represent the non-anomalous cause, the anomalous cause, and the either-non-anomalous-OR-anomalous cause, respectively. Recall that belief assigned to the third element (the truth-set) represents lack of direct evidence for any particular hypothesis; colloquially, ignorance in an expert's knowledge of the truth.

Conditioning of Sensors as Dempster-Shafer Experts

Now that the set of possible hypotheses has been enumerated, BBAs must be developed for each expert. Each expert is treated as a simple BBA, contributing evidence to the hypothesis N for non-anomalous causes of non-detection and contributing the complementary belief mass to the truth-set Θ . Therefore, in general, the BBA for sensor i can be represented as follows: $m_i(N) = p$ and $m_i(\Theta) = 1 - p$

An important aspect of Dempster's conjunctive rule of combination is that the pieces of evidence (the BBAs) should come from independent sources. Since this rule in particular is not idempotent, dependent sources will lead to accounting for the same piece of evidence more than once, leading to skewed belief masses. In this application, most of the sources of evidence can be considered independent without much difficulty; for instance, the proximity of the moon to the space object does not affect background sky brightness or cloud cover. However, some of the particular sensors do share dependencies, which must be handled before creating BBAs.

The approach adopted here is to use a weighted sum to combine the evidence of dependent sources into one expert. In particular, the weather forecast and cloud detection software both provide evidence for sky-cover as a non-detection cause. However, due to the temporal resolution of the All-Sky cloud detection software, it is given more weight in the computation of the BBA. Given an All-Sky sky cover estimate of c_A , a weather forecast sky cover estimate of c_F , and a weighting coefficient of w , the sky cover expert BBA is constructed as:

$$m_C(N) = wc_A + (1 - w)c_F \quad (8)$$

$$m_C(\Theta) = 1 - m_C(N) \quad (9)$$

Re-Tasking Approach

Once the plausibility of an anomaly has been computed, the algorithm must autonomously determine the next course of action. At this point, a decision must be made as to what criterion defines the best course of action. The evidential reasoning approach to anomaly detection yields a new way of viewing follow-on observation tasking in light of the current evidence. Since the gap between belief and plausibility represents a lack of evidence, or ignorance to the truth, a meaningful course of action is to attempt to reduce the belief-plausibility gap as much as possible.

Attempting to find the optimal course of action also requires defining the time-span under consideration. One potential approach is to employ a fixed-horizon scheme, attempting to minimize the

belief-plausibility gap as much as possible over the course of a finite number of steps ahead. One method applicable to this approach is a mixed-integer linear programming approach. However, in this study, we opt for a simpler greedy optimization approach, so-called because it only considers its very next step when determining its next course of action.

In general, the possible courses of action for re-tasking the sensor are: i) search for the maneuvered space object, ii) search for the quiescent (non-maneuvered) space object, iii) continue with the next planned observation, or iv) do nothing. The first option overrides the otherwise next-planned observation searches the space object's reachable volume to confirm the anomalous non-detection hypothesis. The second option also overrides the next-planned observation and opts to search for the next observation in the area predicted by quiescent propagation to confirm the non-anomalous non-detection hypothesis. The third option does not deviate from the next planned observation. The last option opts not to make any observation at the next observation interval.

Since the work in this paper makes the simplifying assumption of a one-object world, the third option listed above is subsumed by option two. Therefore, for this paper, we will formulate the following decision space:

$$\mathcal{D} = \{M, Q, \emptyset\} \quad (10)$$

representing the first, second (and third), and fourth options, respectively, as described above.

The chosen formulation of the hypothesis space as a binary hypothesis pair and their disjunctive combination, as in Eq. (7), simplifies the selection of a re-tasking action. In this situation, maximally reducing ignorance is equivalent to maximally reducing the amount of belief mass attributed to the truth-set, Θ . Reducing the belief-plausibility gap can be formulated as the optimization problem in Eq. (11):

$$\begin{aligned} \min_{d \in \mathcal{D}} \quad J &= \text{ig}(N) = \text{ig}(A) \\ &= m_{F \oplus d}(\Theta) = (m_F \oplus m_d)(\Theta) \end{aligned} \quad (11)$$

where m_F is the fused mass function from the current evidence representing the current knowledge of the system, and m_d is the mass function resulting from the decision $d \in \mathcal{D}$. Since the hypothesis space for this particular application is restricted to 2 elements and their disjoint union, Dempster's rule can be applied to simplify the cost expression to Eq. (12):

$$\begin{aligned} J &= \frac{\sum_{B \cap C = \Theta} m_F(B) m_d(C)}{1 - \sum_{B \cap C = \emptyset} m_F(B) m_d(C)} \\ &= \frac{m_F(\Theta) m_d(\Theta)}{1 - m_F(N) m_d(A) - m_F(A) m_d(N)} \\ &\propto \frac{m_d(\Theta)}{1 - m_F(N) m_d(A) - m_F(A) m_d(N)} \end{aligned} \quad (12)$$

The final simplification notes that the mass function m_F does not change based on the selection of d , so the m_F term in the denominator is simply a scaling factor and can be removed. Also recall that the denominator only measures the level of conflict between the two BBAs. In many cases, this conflict is small and the denominator evaluates to 1. Otherwise, it simply scales the belief mass function as well. This leads to a rather intuitive result: to reduce ignorance in the fused BBA, one should take the action that gathers the least ignorance, minimizing $m_d(\theta)$; in other words, take the action more likely to gather actionable information to confirm a particular hypothesis.

Naturally, since the action d has not yet been executed, m_d is actually an estimate of the mass function that would result from decision d . In particular, for this application, all the experts are represented as simple BBAs: so if $m_d(N) = p$, $m_d(\Theta) = 1 - p$, and $m_d(A) = 0$. Therefore, the algorithm will select the action that has the best estimated chance of confirming any particular hypothesis, since the rest of the belief mass will be assigned to the truth-set as ignorance. This requires methods for estimating the BBAs for each re-tasking action.

Re-Tasking Action BBA Estimation

Determining the optimal re-tasking approach requires estimation of the belief mass associated with each possible re-tasking event to ensure minimization of ignorance. In other words, the telescope should re-task to search for a maneuvered spacecraft if it is more likely to confirm the maneuvered hypothesis than it is to confirm the non-maneuvered hypothesis from a second attempt at observing the area predicted through quiescent dynamics.

When planning a search pattern to establish belief in the anomalous hypothesis, a reachability analysis provides an upper bound on the deviated state, based on known or assumed maximum control authority. In this study, the reachability approach follows a method outlined by Holzinger and Scheeres, particularly Theorem 2, the Reachability Position Maximum theorem.¹⁶ Given an ellipsoidal initial set describing position and velocity state uncertainty bounds and the spacecraft's maximum control authority u_m , the algorithm uses Newton descent to compute the maximal deviation, $\mathbf{d}_{f, mx}$, at time t_f from the state estimate. In the process, it also solves for the initial position and velocity, \mathbf{d}_0 and \mathbf{v}_0 , associated with this maximum deviation, and the Lagrange multiplier λ_0 corresponding to this boundary condition. Therefore, by the Reachability Position Maximum theorem,¹⁶ given the vector $\zeta^T = [\mathbf{d}_0^T, \mathbf{v}_0^T, \mathbf{d}_{f, mx}^T, \lambda_0]$ of decision variables in the descent algorithm, the following system of equations is solved:

$$\begin{bmatrix} \mathbf{d}_{f, mx} \\ \mathbf{d}_{f, mx} \\ \mathbf{0} \end{bmatrix} = \begin{bmatrix} \mathbb{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbb{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbb{I} \end{bmatrix} \phi_z \left(t_f; \begin{bmatrix} \mathbf{d}_0 \\ \mathbf{v}_0 \end{bmatrix}, -2\lambda_0 \mathbf{E} \begin{bmatrix} \mathbf{d}_0 \\ \mathbf{v}_0 \end{bmatrix}, t_0 \right) \quad (13)$$

$$0 = \begin{bmatrix} \mathbf{d}_0 \\ \mathbf{v}_0 \end{bmatrix}^T \mathbf{E} \begin{bmatrix} \mathbf{d}_0 \\ \mathbf{v}_0 \end{bmatrix} - 1 \quad (14)$$

$$\mathbf{E} = \begin{bmatrix} \frac{1}{r_p^2} \mathbb{I} & \mathbf{0} \\ \mathbf{0} & \frac{1}{r_v^2} \mathbb{I} \end{bmatrix} \quad (15)$$

where ϕ_z in Eqn. (13) is the flow function representing the propagation of the state \mathbf{x} and co-state \mathbf{p} trajectory. Equation (14) defines an ellipsoid in \mathbb{R}^6 with shape matrix \mathbf{E} as defined in Eqn. (15), where r_p and r_v are the position and velocity state uncertainties at the initial condition. In this study, 1- σ uncertainty bounds are used in Eqn. (15), though 3- σ or any n - σ bounds could also be used. This system of equations can be solved as 10 equality constraint equations using Newton's method to find the $\zeta \in \mathbb{R}^{10 \times 1}$ that simultaneously satisfies these constraints. Holzinger and Scheeres fully derive this solution procedure, including the required gradients for Newton descent.¹⁶ In this work, the 4D (planar) nonlinear relative dynamics implemented in the reference paper are extended in this application to a full 6D case using similar nonlinear relative dynamics, allowing capture of out-of-plane motion.¹⁷

The reachability computation procedure is illustrated in Fig. 4. Once the maximum position $\mathbf{d}_{f, mx}$ has been calculated, the reachable volume can be determined using Corollary 3 to construct a

maximum bound as a sphere of radius $\|\mathbf{d}_{f,mx}\|$ centered at the expected value of the propagated state estimate.¹⁶ This volume represents the reachable space of the spacecraft at time t_f given the control authority u_m , meaning the spacecraft must reside somewhere within this volume. Without and prior information about the maneuver type or size, each position state within this volume is equally likely, so the principle of indifference can be applied to form a uniform probability density function (PDF) for the spacecraft location. An approximate detection probability is computed by integrating the sensor field-of-view (FOV) over this PDF as depicted in Fig. 4, yielding the probability that the spacecraft is contained within the intersection of the sensor FOV with the position PDF as in Eqn. (16). This is only an approximation because the true reachability set is a subset of this volume.

$$\mathbb{P}_D = \frac{\int \int_{L_m} f_R(\alpha, \delta) d\alpha d\delta}{\int \int_{\mathcal{R} \setminus \{L_k\}_{k=1}^{m-1}} f_R(\alpha, \delta) d\alpha d\delta} \quad (16)$$




Figure 4. Illustration of reachable space bounded by the maximum reachable distance, $\mathbf{d}_{f,mx}$, projected onto image plane and in 3D, including intersection of sensor FOV with reachable volume

If the telescope is commanded to re-task for detection in this reachable volume, the optimal sensor tasking maximizes this detection probability, which yields the next sensor pointing command. It is important to note that already-observed areas of the reachability volume should be neglected from the total PDF since the spacecraft was not found. Therefore at the m^{th} observation attempt, the portion of the reachable volume associated with all previous unsuccessful attempts $\{L_k\}_{k=1}^{m-1}$ should be neglected when computing the PDF. In contrast, the non-maneuver decision belief mass is simply computed in the same way as the original observation. The state estimate is propagated from t_0 to t_f to yield the next pointing vector.

However, the estimated decision event belief mass is not based solely on the reachability detection probability, but also on the predicted probability of detection or non-detection based on the other factors already discussed (e.g. cloud cover). The anomalous and non-anomalous search hypotheses both provide a next pointing vector, so these can be used to determine estimated BBAs for sky brightness, cloud cover, moon proximity, and any other non-detection phenomena. These belief

functions are all fused, as discussed previously for expert belief mass fusion, to form the final BBA for each course of action.

IMPLEMENTATION

This section contains details about the software and hardware implementation of the loss-of-custody algorithm. For this study, the set of experts available was reduced to match the available sensors and software at the Georgia Tech observatory. An All-Sky camera and weather forecast parser are used for cloud detection, and a sky brightness sensor measures background sky brightness.

All-Sky Cloud Detection

The SBIG All-Sky 340 camera, shown in Fig. 5, is a relatively low-resolution monochrome CCD with a fisheye lens that allows it to achieve a horizon-to-horizon field of view. Relevant parameters for the All-Sky camera are listed in Table 1.

Table 1. SBIG All-Sky 340 parameters

Parameter	Value	Units
Field of View	185	degrees
Resolution (H x V)	640 x 480	pixels
Focal Length	1.4	millimeters
Focal Ratio	f/1.4	-



Figure 5. SBIG All-Sky 340

Optical sensors are notoriously bad at detecting clouds using pixel-to-pixel derivative-based methods because clouds often exhibit wide brightness variations due to internal structure. Additionally, while bright areas in the clouds might be detected, the wispy cloud portions, which still very-much affect seeing, are often missed in traditional pixel-wise derivative-based blob detection algorithms. Other methods proposed to handle cloud detection utilize the difference between the color channels to detect clouds against the blue sky.¹⁸ However, since the All-Sky camera available is monochrome and the relevant imaging time is night, this method cannot be applied here.

Therefore, this paper implements an optical-flow-based object detection algorithm to utilize pixel-to-pixel and frame-to-frame derivatives for computing motion between frames. The optical flow algorithm assumes brightness-constancy of objects and computes a magnitude of motion from frame-to-frame, performing best when the displacement from frame-to-frame is relatively low. These characteristics are well-suited to cloud detection.¹⁹ For the observatory roof-mounted All-Sky camera, the only objects moving in the frame are clouds, the lighting conditions do not change considerably, and the 30-seconds between subsequent frames does not allow clouds to drift too far. Additionally, the computation of frame-to-frame motion allows the ability to predict sky cover in other areas of the sky that are of interest for the next observation.

Before attempting optical flow, though, the barrel distortion of the fisheye lens must be corrected back to a rectilinear (pinhole projection) mapping. The lens in the SBIG All-Sky 340 exhibits *f-theta* distortion, which is a function of the focal length of the lens and the object's distance from the boresight of the camera.²⁰ Therefore, each image from the SBIG camera is undistorted as

shown in Fig. 6, which results in linear motion being correctly represented from frame-to-frame, a useful result for predicting sky cover from pixel velocity estimates. The conversion does cut off the edges of the distorted image, but this does not hurt performance, particularly in an urban environment like Atlanta. The sky near the skyline is much brighter, not ideal for observation, and the mounting limitations of the telescope do not allow observation that close to the horizon. The outlined conversion area, shown in Fig. 6, was chosen to be just above the tallest building in the Atlanta skyline, which still captures the majority of the sky, particularly the usable area for observations. Although the images in Fig. 6 appear similar sized in print, the undistorted images are actually 2.5 times larger in each dimension.




Figure 6. All-Sky cloud detection

The Horn-Schunck optical flow algorithm formulates the energy function in Eq. (17) subject to a candidate flow field (u, v) , seeking to minimize energy by modifying the flow field through a gradient descent.²¹ Pixel-wise gradients (I_x and I_y) are computed using central differencing and frame-to-frame derivatives (I_t) are computed using backward differencing. The time difference between the frames is known based on the timestamp in the image metadata (nominally 30 seconds between frames). After a user-defined convergence criterion is met, the magnitudes of the flow field (u, v) can be evaluated to determine which pixels contain cloud and which contain empty sky.

$$E(x, y, t) = \int \int ([I_x(x, y, t)u(x, y, t) + I_y(x, y, t)v(x, y, t) + I_t(x, y, t)]^2 + \|\nabla u(x, y, t)\|^2 + \|\nabla v(x, y, t)\|^2) dx dy \quad (17)$$

The optical flow algorithm excels at detecting the fainter, wispy portions of the clouds, but does not detect the flat, bright areas of heavy cloud with either pixel-wise or frame-wise derivatives. Therefore, the optical flow algorithm is augmented with a component that checks for bright pixels above a threshold. If a particular pixel exceeds either the brightness or flow velocity magnitude thresholds, it is deemed a cloud. Figure 6 shows the result for sample images. In this study, the thresholds for brightness and flow velocity magnitude are $\frac{75}{255}$ and $1.5e - 4$ respectively, and the optical flow PDF is evolved at an artificial timestep of $1ms$.

Further improvement could be made on this algorithm by incorporating a probabilistic threshold. However, for this application, a simpler approach of using a 3×3 Gaussian kernel to smear the edges of the cloud map was implemented to include a probabilistic component.

Weather Forecasts

Weather forecasts are used to augment information from the All-Sky cloud detection algorithm. A python script parses information from the National Oceanic and Atmospheric Administration website*, retrieving weather forecasts for data such as sky cover, humidity, and precipitation potential. The NOAA website uses the NWS National Digital Forecast Database (NDFD), a model which compiles information on numerous statistics, including sky cover, temperature, and humidity. In future studies, the authors wish to use more weather statistics to predict other atmospheric properties than just low cloud cover. In particular, forecasts on high clouds that are difficult to detect even with optical flow could further inform seeing conditions. A sample weather forecast is shown in Fig. 7.

GT-SORT Time (EDT)	Weather	Forecast Sky Cover	Update: Jan 06, 2016 - 18:22 Precip?	Temp (F)	Dew Pt (F)	Rel Hum
01/06 18:00		81%	0%	42	18	37
01/06 19:00		81%	0%	41	19	41
01/06 20:00		79%	0%	40	20	45
01/06 21:00		78%	0%	39	22	50
01/06 22:00		77%	0%	38	23	53
01/06 23:00		77%	0%	38	23	54
01/07 0:00		76%	0%	37	23	55
01/07 1:00		75%	0%	37	23	56
01/07 2:00		75%	0%	36	23	60
01/07 3:00		76%	0%	36	24	63
01/07 4:00		76%	0%	35	25	68
01/07 5:00		77%	0%	34	26	71
01/07 6:00		77%	0%	34	26	74
01/07 7:00		78%	0%	34	27	75
01/07 8:00		78%	0%	33	28	82
01/07 9:00		78%	0%	35	30	81
01/07 10:00		78%	0%	38	32	77
01/07 11:00		78%	0%	42	34	73
01/07 12:00		78%	0%	45	35	70
01/07 13:00		78%	0%	50	37	61
01/07 14:00		78%	0%	54	38	55
01/07 15:00		79%	0%	54	39	56
01/07 16:00		80%	1%	54	40	58
01/07 17:00		81%	1%	53	40	61

Figure 7. Sample weather forecast output

Sky Brightness

Assessing the effect of background sky brightness on the detection of a space object requires knowledge of numerous properties of the sensors, the spacecraft, and the environment. The goal is to develop a probability of detection based on the observed sky brightness from the Unihedron Sky Quality Meter, which measures sky brightness in units of visual apparent magnitudes per square-arcsecond. The development of a detection probability based on required algorithm SNR, sky brightness, space object brightness, and other optical properties is presented by Coder and Holzinger.²² Random variables associated with the number of incident photons are represented through Poisson distributions, and the central limit theorem is applied to approximate the number of photons as a Gaussian distribution. Statistics are then developed based on optical and environmental properties to allow computation of the probability that the signal is greater than the noise using a Gaussian cumulative distribution function (CDF). The particular equation used for probability of detection, from Coder and Holzinger,²² is shown in Eq. (18).



Figure 8. Unihedron Sky Quality Monitor

$$\mathbb{P}[\Gamma_{SO} > \text{SNR}_{alg}\sigma_n] = \frac{1}{2} \left[1 - \text{erf} \left(\frac{\text{SNR}_{alg}\sigma_n - \mu_{SO}}{\sqrt{2}\sigma_{SO}} \right) \right] \quad (18)$$

*NOAA National Weather Service Forecast Office, Hourly Weather Forecast, Atlanta, GA: <http://forecast.weather.gov/MapClick.php?lat=33.7629&lon=-84.4226&unit=0&lg=english&FcstType=graphical>

Telescope and optical system parameters are taken from the Georgia Tech Space Object Research Telescope (GT-SORT), a Raven-class telescope installed at the Georgia Institute of Technology on-campus observatory.^{22,23} Selected environmental properties for observation from midtown Atlanta are also listed in Table 2. Coder and Holzinger provide a thorough discussion of each of these parameters and their origin.²² Many of these parameters are constants, fixed by the chosen hardware. Others, such as the atmospheric transmittance, are estimates based on the observation environment. More accurate estimates of atmospheric transmittance could be obtained by comparing the expected brightness of stars in the FOV to their actual observed values.

Parameter	Value	Units
Field of View (H x V)	14.2 x 11.4	arc-minutes
Resolution (H x V)	2736 x 2192	pixels
Focal Ratio	f/6	-
Aperture Diameter	0.5	m
Quantum Efficiency	0.74	-
Atmospheric Transmittance	0.50	-
Optical Transmittance	0.90	-
Secondary Transmittance	0.84	-
Zero-Magnitude Irradiance	5.6e10	photons/s/m ²
Required Algorithm SNR	4	-

Table 2. GT-SORT telescope, environment, and algorithm parameters

This detection scheme is conditioned on an expected brightness magnitude for the space object, listed in tables in each test case 2. Photometric modeling of spacecraft is an active area of research and can be difficult to perform accurately due to attitude maneuvering of spacecraft or unknown physical parameters. For simplicity in this study, the value chosen is selected such that, for the nominal sky brightness in each simulated case, the detection probability is not exactly 0 or 1.

EXPERIMENTAL RESULTS

Data from each of the listed sensors was collected over a number of nights to perform an analysis of the non-detection hypothesis testing and re-tasking decision making algorithms. The target object selected for this study is Echostar 11 (COSPAR designation: 2008-035A), a geostationary Earth orbiting (GEO) satellite positioned at 110 degrees West longitude. It has a radar large cross section and is positioned over the United States which makes it ideal for observation from the Georgia Tech observatory. Figure 9 shows a recent 10-second unfiltered exposure of Echostar 11 taken using GT-SORT. The result shown has been dark-frame subtracted and the colors inverted for print. The faint dot near the center of the GT-SORT image is the tracked satellite, while a number of stars can also be seen streaking through the background. Similar recoloring is performed for the accompanying All-Sky camera, only a few light clouds can be seen in this image. The cyan dot indicates the position of Echostar 11.

The test cases elaborated below utilize real data taken from the evening of January 10, 2016. This night included times of moderate cloud cover as well as times of clear skies. In these test cases, the cloud cover and sky brightness data are processed under an assumed observation attempt of the geostationary satellite, Echostar 11, with assumed parameters listed in Table 3. In each case, the spacecraft is assumed not to be detected, requiring the use of the non-detection hypothesis




Figure 9. Echostar 11 images taken using GT-SORT and the All-Sky camera on January 14, 2016 at 3:35 (UTC)

testing algorithm to determine whether detection was prevented by anomalous or non-anomalous causes. The reachability algorithm is then applied in re-tasking the satellite to minimize ignorance in the next observation. Due to the satellite’s GEO orbit, determining predicted cloud cover or sky brightness at its next location is simplified since it appears stationary with respect to the observer.

Parameter	Value
Target Longitude	-110 deg
Previous Observation Date (UTC)	Jan 10, 2016
Previous Observation Time (UTC)	4:30
Control Authority ^a	$1 \times 10^{-7} \frac{m}{s^2}$
Visual Magnitude	8.974
Projected FOV Width @ GEO	209 km
Projected FOV Height @ GEO	168 km
Cloud Expert Weight (w)	0.8

Table 3. Assumed target (Echostar 11) operational and previous estimate parameters

^aSpecific information on the control authority was not available. The value chosen sizes the reachability volume such that it exceeds the projected field of view.

Test Case 1: Clear, Dark Skies

The first test case involves non-detection when there are no obvious non-anomalous contributing factors. Observation conditions for the first test case are enumerated in Table 4. In this test case, the weather forecast predicts low cloud coverage at this particular hour. The All-Sky cloud detection algorithm determines that there is no cloud cover in the observation location. Additionally, at the time of this observation, the background sky brightness is not significantly higher than normal, meaning the spacecraft should be mostly detectable.

The observed values are converted to BBAs as shown in Table 5. Recall that the cloud cover expert BBA, m_C , is created through a weighted combination of the weather forecast and All-Sky

Table 4. Observation conditions for Test Case 1

	Parameter	Value
Measured	Observation Date (UTC)	1/11/16
	Observation Time (UTC)	4:21:27
	Integration Time (s)	2.7913
	Forecast Cloud Cover	10 %
	All-Sky Cloud Cover	0 %
	Sky Irradiance ($\frac{m_v}{\text{arcsec}^2}$)	17.41
Predicted	Forecast Cloud Cover	10 %
	All-Sky Cloud Cover	0 %
	Sky Irradiance ($\frac{m_v}{\text{arcsec}^2}$)	17.41
	Max Reachable Distance	243 km
	Detection Probability (M)	23 %

Table 5. BBAs for Test Case 1

BBA	$\{N\}$	$\{A\}$	Θ
m_C	0.020	–	0.980
m_B	0.022	–	0.978
$m_{C\oplus B}$	0.042	0	0.958
m_{C_Q}	0.020	–	0.980
m_{B_Q}	0.022	–	0.978
$m_{C_Q\oplus B_Q}$	0.042	0	0.958
$m_{C\oplus B\oplus C_Q\oplus B_Q}$	0.081	0.096	0.919
m_{C_M}	0.020	–	0.980
m_{B_M}	0.022	–	0.978
m_{D_M}	–	0.230	0.770
$m_{C_M\oplus B_M\oplus D_M}$	0.032	0.223	0.745
$m_{C\oplus B\oplus C_M\oplus B_M\oplus D_M}$	0.064	0.215	0.721

cloud cover. In this case, most of the belief mass for both expert BBAs is attributed to Θ , the truth-set, representing the fact that neither expert could conclusively say anything about the non-detection aside from the fact that their area-of-expertise is highly unlikely to be the cause. Therefore, the fused BBA $m_{C\oplus B}$ is shown to also have most of its belief mass attributed to Θ .

Next, the BBAs associated with the decision to check the quiescent hypothesis and maneuver hypothesis are shown, denoted with Q and M subscripts respectively. Since the observation cadence is low and the line-of-sight to the space object hasn't changed, the local brightness conditions are not expected to change considerably. Additionally, the propagated cloud cover estimate still shows clear skies at the observation location during the next observation attempt. Finally, note that the target spacecraft in these simulations is in GEO, so the cloud cover and sky brightness estimates for both decisions are identical. This is not a requirement; in general, the maneuvered and quiescent decisions can refer to different areas of the sky and therefore obtain different belief masses from individual experts.

The maneuver-detection decision BBA contains one extra term, associated with the reachability results. This refers to the estimated belief that an anomaly search will be successful in detecting the space object given the field of view and the size of the reachability space. The added term also finally provides belief mass to the anomaly hypothesis, since a successful detection during the anomaly search is the only expert in this scheme that directly provides evidence of an anomaly.

Decision	BBA	bel(N)	pl(N)	bel(A)	pl(A)
Observation ($d = \emptyset$)	$m_{C\oplus B}$	0.042	1.0	0.0	0.958
Quiescent ($d = Q$)	$m_{C\oplus B\oplus C_Q\oplus B_Q}$	0.081	1.0	0.0	0.919
Maneuvered ($d = M$)	$m_{C\oplus B\oplus C_M\oplus B_M}$	0.064	0.785	0.215	0.936

Table 6. Belief and Plausibility for Test Case 1

In this first test case, all experts are indicating that the spacecraft should have been easily observable, leading a high anomaly plausibility after the initial observation attempt. The re-tasking

algorithm determines it is advantageous to search for the missing space object since the detection probability is high, as evidenced by the decreased level of ignorance in the fused BBA.

In cases similar to test case 1, the algorithm will tend to opt to begin searching for a maneuver since the observation conditions are considered otherwise pristine. However, if the observation conditions in the maneuver search area are worse in comparison to the quiescent observation area conditions (e.g. more clouds or brighter skies locally), the fused belief mass associated with anomaly will be less and the algorithm could choose to avoid searching in a poor observation area.

Test Case 2: Localized Cloud Cover

The second test case involves non-detection in the presence of local cloud cover. Observation conditions for the second test case are enumerated in Table 7. In this test case, the weather forecast predicts low cloud coverage at this particular hour. However, the All-Sky cloud detection algorithm computes a high probability of cloud cover at the observation location during the original attempt, followed by clearer skies at the next observation attempt. Similar to the previous test case, the background sky brightness is not significantly higher than normal, so sky brightness does not contribute significant belief for non-anomaly.

Table 7. Observation conditions for Test Case 2

	Parameter	Value
Measured	Observation Date (UTC)	1/11/16
	Observation Time (UTC)	4:30:01
	Integration Time (s)	2.7913
	Forecast Cloud Cover	10 %
	All-Sky Cloud Cover	71.4 %
	Sky Irradiance ($\frac{m_v}{\text{arcsec}^2}$)	17.41
Predicted	Forecast Cloud Cover	10 %
	All-Sky Cloud Cover	8.2 %
	Sky Irradiance ($\frac{m_v}{\text{arcsec}^2}$)	17.41
	Max Reachable Distance	243 km
	Detection Probability (M)	23 %

Table 8. Basic belief assignments for Test Case 2

BBA	$\{N\}$	$\{A\}$	Θ
m_C	0.592	–	0.408
m_B	0.022	–	0.978
$m_{C\oplus B}$	0.601	0	0.399
m_{C_Q}	0.088	–	0.912
m_{B_Q}	0.022	–	0.978
$m_{C_Q\oplus B_Q}$	0.108	0	0.892
$m_{C\oplus B\oplus C_Q\oplus B_Q}$	0.644	0	0.356
m_{C_M}	0.088	–	0.912
m_{B_M}	0.022	–	0.978
m_{D_M}	–	0.230	0.770
$m_{C_M\oplus B_M\oplus D_M}$	0.085	0.210	0.704
$m_{C\oplus B\oplus C_M\oplus B_M\oplus D_M}$	0.582	0.096	0.321

The observed values above are converted to BBAs as shown in Table 8. Since the local cloud cover is moderately high, belief mass for the cloud cover expert in this case is more evenly split between N and Θ . Therefore, the fused BBA $m_{C\oplus B}$ also splits its belief mass between the non-anomalous hypothesis and the truth-set Θ . The BBAs associated with the decision to check the quiescent hypothesis and maneuver hypothesis are also shown in Table 8. In this case, the local cloud cover from the initial observation has passed in time for the second attempt, meaning predicted belief in the non-anomalous hypothesis is much lower. Since the time since the previous estimate is still roughly 24 hours, the reachability space size has not changed substantially, so the probability of maneuver detection has not changed either.

In the second test case, the algorithm once again determines that searching for a maneuver is more advantageous than simply searching again along the quiescent trajectory, but the margin is slimmer

Decision	BBA	bel(N)	pl(N)	bel(A)	pl(A)
Observation ($d = \emptyset$)	$m_{C \oplus B}$	0.601	1.0	0.0	0.399
Quiescent ($d = Q$)	$m_{C \oplus B \oplus C_Q \oplus B_Q}$	0.644	1.0	0.0	0.356
Maneuvered ($d = M$)	$m_{C \oplus B \oplus C_M \oplus B_M}$	0.582	0.904	0.096	0.418

Table 9. Belief and Plausibility for Test Case 2

than in the first test case. This can be seen in the smaller post-decision truth-set belief masses in case two. The estimated ignorance after each decision is similar, slightly favoring a search for maneuver due to its predicted ability to provide direct evidence for maneuver.

It is worth noting that, since the quiescent and maneuvered decision scenarios share all but one contributing expert in this scenario, the quiescent hypothesis cannot contribute any less ignorance than the maneuver search hypothesis: its ignorance will only be equal to or greater than the maneuvered hypothesis. This is not true in general, but this effect is accentuated in these test cases because the propagated uncertainty and reachability volumes are coincident to the resolution of the non-telescope sensors.

CONCLUSIONS

This study applied modern data fusion and autonomous decision-making processes to an SDA scenario. The JDL/DFIG framework allowed a systematic loss-of-custody algorithm to be developed and implemented. The Dempster-Shafer evidential reasoning approach allowed for better decision analyst modeling, providing a robust framework for the fusion of many sensors or experts in different hypotheses. Methods for conditioning these sensors as Dempster-Shafer experts were developed and tested using real-world data from the Georgia Tech Observatory. Additionally, a re-tasking algorithm based on reducing the gap between belief and plausability was developed. The test cases shown illustrate the algorithm’s ability to deduce the cause of missed-detection even with few experts, re-tasked itself to look for a maneuvered spacecraft based on the lack of a clear cause for missed-detection. Preliminary results are promising, but further testing on a wider variety of SSA experts and spacecraft scenarios will lend greater insight into the applicability of this evidential-reasoning-based algorithm to autonomous SDA decision-making.

FUTURE WORK

There are a number of areas for improvement in this algorithm, as well as assumptions that should be relaxed in future study. For instance, the sky brightness estimation can be greatly improved by the addition of the Garstang model, which estimates sky brightness at a given zenith angle based on nearby light-pollution sources and atmospheric conditions such as aerosol density.²⁴ While this paper provides a framework for application of Dempster-Shafer experts, more non-anomaly experts must be added in the future to better cover possible non-anomalous causes. A combined moon brightness and ephemeris model can be applied to predict obstructions by the moon. Perhaps the most important area of future work, though is to relax the one-object world assumption, which increases the complexity of the decision space. Test cases that do not involve solely GEO satellites will be investigated as well to better exercise the re-tasking decision.

ACKNOWLEDGMENTS

The authors would like to thank Dr. Kim Luu for inspiring this work, as well as her guidance and operational insights during its development. This material is based upon work supported by the

REFERENCES

- [1] J.-C. Liou, "Modeling the Large and Small Orbital Debris Populations for Environment Remediation," tech. rep., NASA Orbital Debris Program Office, Johnson Space Center, Houston, TX, June 2014.
- [2] A. P. Dempster, "The Dempster-Shafer Calculus for Statisticians," *International Journal of Approximate Reasoning*, 2007.
- [3] Azimirad and Haddadnia, "The Comprehensive Review on JDL Model in Data Fusion Networks: Techniques and Methods," *International Journal of Computer Science and Information Security*, Vol. 13, January 2015.
- [4] Steinberg, Bowman, and White, "Revisions to the JDL Data Fusion Model," *AeroSense'99*, International Society for Optics and Photonics, March 1999, pp. 430–441.
- [5] W. F. Caselton and W. Luo, "Decision Making with Imprecise Probabilities: Dempster-Shafer Theory and Application," *Water Resources Journal*, Vol. 28, December 1992, pp. 3071–3083.
- [6] W. L. Oberkampf, J. C. Helton, and K. Sentz, "Mathematical Representation of Uncertainty," *Non-Deterministic Approaches Forum*, Seattle, WA, April 2001, pp. 1–23.
- [7] R. R. Yager, "Arithmetic and Other Operations on Dempster-Shafer Structures," *International Journal of Man-Machine Studies*, Vol. 25, 1986, pp. 357–366.
- [8] R. R. Yager, "On the Dempster-Shafer Framework and New Combination Rules," *Information Sciences*, Vol. 41, 1987, pp. 93–137.
- [9] L. A. Zadeh, "A Simple View of the Dempster-Shafer Theory of Evidence and its Implication for the Rule of Combination," *AI Magazine*, Vol. 7, No. 2, 1986, pp. 85–90.
- [10] R. Haenni, "Shedding New Light on Zadeh's Criticism of Dempster's Rule of Combination," *8th International Conference on Information Fusion*, Vol. 2, IEEE, July 2005.
- [11] P. Smets and R. Kennes, "The Transferable Belief Model," *Artificial Intelligence*, Vol. 66, 1994, pp. 191–234.
- [12] T. Denoex, "Conjunctive and Disjunctive Combination of Belief Functions Induced by Nondistinct Bodies of Evidence," *Artificial Intelligence*, Vol. 172, 2008, pp. 234–264.
- [13] R. R. Yager and L. Liu, *Classic Works of the Dempster-Shafer Theory of Belief Functions*. Springer, 2008.
- [14] A. D. Jaunzemis, M. V. Mathew, and M. J. Holzinger, "Control Metric Maneuver Detection with Gaussian Mixtures and Real Data," *25th AAS/AIAA Spaceflight Mechanics Conference*, Williamsburg, VA, January 2015.
- [15] A. D. Jaunzemis, M. V. Mathew, and M. J. Holzinger, "Control Cost and Mahalanobis Distance Binary Hypothesis Testing for Spacecraft Maneuver Detection," *Journal of Guidance, Control, and Dynamics*, 2015 (submitted).
- [16] M. J. Holzinger and D. J. Scheeres, "Reachability Results for Nonlinear Systems with Ellipsoidal Initial Sets," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 48, April 2012, pp. 1583–1600.
- [17] C. W. T. Roscoe, "Reconfiguration and Recovery of Formation Flying Spacecraft in Eccentric Orbits," Master's thesis, University of Toronto, 2009.
- [18] A. Kazantzidis, P. Tzoumanikas, A. F. Bais, S. Fotopoulos, and G. Economou, "Cloud Detection and Classification with the Use of Whole-Sky Ground-Based Images," *Atmospheric Research*, Vol. 113, September 2012, pp. 80–88.
- [19] Solar Thermal Group, Australian National University, *Cloud Tracking with Optical Flow for Short-Term Solar Forecasting*, Canberra, Australia, 2012.
- [20] J. Courbon, Y. Mezzour, L. Eck, and P. Martinet, "A Generic Fisheye Camera Model for Robotic Applications," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2007, pp. 1683–1688.
- [21] B. K. Horn and B. G. Schunck, "Determining Optical Flow," *Techniques and Applications of Image Understanding* (J. J. Pearson, ed.), Vol. 281, Washington, D.C., April 1981.
- [22] R. Coder and M. J. Holzinger, "Multi-Objective Design of Optical Systems for Space Situational Awareness," *Acta Astronautica*, 2015 (submitted).
- [23] R. D. Coder and M. J. Holzinger, "Sizing of a Raven-class Telescope using Performance Sensitivities," *Advanced Maui Optical and Space Surveillance Technologies Conference*, Wailea, HI, September 2013.
- [24] J. A. J. Birriel, "A Simple, Portable Apparatus to Measure Night Sky Brightness at Various Zenith Angles," *Journal of the American Association of Variable Star Observer*, Vol. 38, No. 221-229, 2010.