

1st IAA Conference on Space Situational Awareness (ICSSA)

Orlando, FL, USA

IAA-ICSSA-17-0X-XX

Judicial Evidential Reasoning for Decision Support Applied to Orbit Insertion Failure

Andris D. Jaunzemis^a, Dev Minotra^b, Marcus J. Holzinger^c, Karen M. Feigh^d,
Moses W. Chan^e, Prakash P. Shenoy^f

^aNSF Fellow, Georgia Institute of Technology, Atlanta, GA,
adjaunzemis@gatech.edu

^bPost-Doctoral Fellow, Georgia Institute of Technology, Atlanta, GA,
dev.minotra@aerospace.gatech.edu

^cAssistant Professor, Georgia Institute of Technology, Atlanta, GA,
holzinger@ae.gatech.edu

^dAssociate Professor, Georgia Institute of Technology, Atlanta, GA,
karen.feigh@gatech.edu

^eTechnical Fellow, Lockheed Martin Corporation, Sunnyvale, CA,
moses.w.chan@lmco.com

^fDistinguished Professor of Artificial Intelligence, University of Kansas, Lawrence, KS,
pshenoy@ku.edu

Keywords: sensor network, situational awareness, decision support systems, Dempster-Shafer theory, hypothesis resolution

Realistic decision-making often occurs with insufficient time to gather all possible evidence before a decision must be rendered, requiring an efficient process for prioritizing between potential action sequences. This work aims to develop a decision support system for tasking sensor networks to gather evidence to resolve hypotheses in the face of ambiguous, incomplete, and uncertain evidence. Studies have shown that decision-makers demonstrate several biases in decisions involving probability judgment, so decision-makers must be confident that the evidence-based hypothesis resolution is strong and impartial before declaring an anomaly or reacting to a conjunction analysis. Providing decision-makers with the ability to estimate uncertainty and ambiguity in knowledge has been shown to augment effectiveness. The proposed framework, judicial evidential reasoning (JER), frames decision-maker questions as rigorously testable hypotheses and employs an alternating-agent minimax optimization on belief in the null proposition. This approach values impartiality in addition to time-efficiency: an ideal action sequence gathers evidence to quickly resolve hypotheses while guarding against bias. JER applies the Dempster-Shafer theory of belief functions to model knowledge about hypotheses and quantify ambiguity, and adversarial

optimization techniques are used to make many-hypothesis resolution computationally tractable. This work includes derivation and application of the JER formulation to a GTO insertion maneuver anomaly scenario.

1. Introduction

Endsley defines situation awareness as “the perception of the elements in the environment within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future” [1]. Space situational awareness (SSA) is particularly 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. While Endsley’s construct of situation awareness refers to a cognitive state in an individual’s mind, space situational awareness is used as an organizational construct referring to the distributed state of knowledge at the organizational level. Potential SSA needs include maintaining catalogs of space object state observations [12, 13], detecting maneuvers or other anomalies [14], and estimating control modes or behavior [15, 4]. Currently, there are over 20,000 trackable objects in the space object catalog [2, 3] ranging from decommissioned rocket bodies to active telecommunications assets to university science and technology experiments. While Earth orbit is a vast volume, useful or strategic orbit regimes (e.g. low Earth orbit (LEO), Geostationary Earth Orbit (GEO), sun-synchronous LEO) have quickly become congested and contested [4]. The number of trackable space objects is continually growing with expanded use of small spacecraft technologies [6] and increased sensor capabilities. Growing clutter poses safety concerns, accentuated by the high-profile LEO collision event in 2009 between a defunct COSMOS satellite and an active Iridium satellite [7]. With such diverse involvement in the space arena, there is a large economic incentive to understand the space environment to ensure continued operation of assets.

Maintaining SSA is essential to the command and control missions of the Joint Space Operations Center (JSpOC) [5]. Discourse and activity in SSA increasingly focuses on decision-making in the presence of limited resources, uncertain information, and a contested space environment. Establishing protocols and regulations in the use of space depends upon the “availability of quantifiable and timely information regarding the behavior of resident space objects” [4]. Unfortunately, constraints imposed by non-linear orbital dynamics and the disparity between the number of space objects and the number of sensors hinder the ability to reliably provide information on maneuvers or other events. An increasing 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 [8, 5].

Tracking techniques used in the space surveillance system still largely rely upon models and applications from the 1950s and 1960s [9], which are human-intensive. For instance, current space object custody tasking requires human analyst to compare candidate tasking schedules while incorporating constraints such as observation conditions (e.g. sky brightness, cloud cover). In the event that an object is not detected, a human analyst may be required to inspect the observation conditions visually before declaring lost custody or anomaly. This approach is reactive and rigid, necessitating improved automated approaches to data collection and processing that incorporate auxiliary sensor data to operate in a more predictive manner and dynamically adjust

the algorithm objectives and actions. As the space object population increases, the amount of data required to maintain SSA also increases [10], which makes human-in-the-loop involvement in space surveillance particularly troublesome and motivates the development of autonomous sensor tasking capabilities.

The sensor tasking or sensor scheduling problem addresses how to obtain, process, and utilize information about the state of the environment [11]. The SSA sensor tasking problem is a high-dimensional, multi-objective, mixed-integer, non-linear optimization problem, so current approaches focus on tractable sub-problems (e.g. single objectives, limited target objects, limited sensors). For instance, catalog maintenance requires observations of many different space objects. Entropy-minimizing or information-maximizing methods, as characterized through covariance estimates, minimize state estimate uncertainty for all catalog objects [16, 12]. Other objectives may require more data of specific targets or events. Space object association may be handled by quantifying a state anomaly or maneuver required to associate two uncorrelated tracks (UCTs) [17, 14], classification methods may employ taxonomies trained on representative space object feature sets to categorize space objects [18], and attitude or control mode estimation requires many observations of a single object to develop a light curve, a time-history of photons received from the target space object [15]. The competing objectives are generally not complementary, especially given limited sensor resources, so the different objectives prefer different tasking approaches.

Sensor information must be fused into a coherent understanding of the environment via association, correlation, and combination [11]. In classical Bayesian approaches, sensor data is used to form deterministic probabilities 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 represented in a strictly binary manner, needing to account for uncertain information and 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 ambiguity, leading to more realistic modeling of human analyst processes [19, 20, 21]. Dempster-Shafer theory has gained significant traction in various applications, including classification [22, 23], monitoring and fault detection [24, 25], and decision-making [26].

This work formulates a sensor tasking algorithm to eliminate both conflict and ambiguity in user-prioritized specific hypotheses, gathering evidence that yields a more precise understanding of the relevant hypotheses. Beginning with an analysis of the work domain of an SSA decision-maker, the three primary concerns that the proposed framework addresses are discussed: hypothesis abstraction, ambiguity aversion, and unbiased hypothesis resolution. The proposed framework is developed and simulated results are presented for an orbit insertion maneuver failure.

2. Theory

2.1. Work-Domain Analysis

The Joint Space Operations Center (JSpOC) plays the role of maintaining space situational awareness to carry out command and control operations pertaining to the protection of civil, defense, and national space assets [27]; tasks in JSpOC support the monitoring, integration, and reporting of space asset state changes [5]. In order

to respond to routine and anomalous events in space and maintain the space object catalog, JSpOC is responsible for issuing requests to task sensors maintained by other organizations such as the US Space Surveillance Network. Various tasks compete for the limited number of sensors available, and various human analysts are involved in making decisions pertaining to sensor tasking.

To support decisions in SSA processes such as sensor tasking and conjunction assessment, decision-support systems should provide the analysts with the relevant information in such a manner as to help the operators identify possible conjunctions and take the correct action to avoid likely conjunctions as soon as possible while minimizing unnecessary maneuvering.

2.1.1. What is a work-domain analysis?

A work-domain analysis [28] is a method to obtain requirements for support systems through the development of a work-domain model that describes work that comprises a socio-technical system. A work-domain model is a complete description of a system in the form of multi-level diagrams referred to as abstraction hierarchies (AH). In an abstraction hierarchy, each level describes the system at a different level of abstraction. The highest level describes the purpose of the work domain and the lowest levels list components and their attributes. The levels are connected to each other via a why-what-how decomposition. Elements above the current level describe why an element in the current level is included; elements below the current level describe how an element will be achieved. Burns and Hajdukiewicz [29] provide an example of a work-domain analysis of a car, which has a purpose of transporting people from A to B safely (top level: purpose). This is achieved through conservation of energy and energy flow (second level: principles, priorities, and values), which is achieved with the help of combustion (third level: processes). Combustion requires a combustion chamber (fourth level: components), which has a particular location and color (fifth level: attributes). Here, we have listed only one element from each level of this example provided by Burns and Hajdukiewicz (2004). This “means-ends” structure to organize information on displays is intended to support decision making in a wide range of scenarios.

2.1.2. Why is a work-domain analysis being conducted?

The value of a work-domain analysis can be seen in displays that organize information in a hierarchical structure in accordance with the abstraction hierarchy. Vicente states that prior studies applying work-domain analysis have shown that the use of functional level information (level 2 and level 3 of an abstraction hierarchy) can be critical in task performance [30]. Solving a problem can require operators to look at information and variables typically present in the lower levels (level 4 and level 5); however, initiation of the problem-solving process with the functional levels (levels 2 and 3) and using means-ends linkages between higher and lower levels is more efficient as it constrains the operators goal-directed search within information relevant to the given problem [31]. Vicente has argued that the an abstraction hierarchy can represent problem spaces in various tasks as this has been shown in empirical studies that map abstraction hierarchy representations to verbal protocols or behaviors of expert performers [31]. These mappings have been verified in studies of expert computer programmers, world-class chess players, nuclear power plant operators, and equipment repair professionals [31].

Table 1: Abstraction hierarchy for SSA decision-support

Levels	Abstraction Hierarchy Elements					
Purpose	Space Asset Safety and Security	National Security				
Principles, Priorities, Values	Space Object Motion - Physical Laws and Orbital Mechanics (principles)	Sensing Capability - Physical Constraints (principles)	Critical Assets Should be Given Priority (priorities)	Asset State Information Continuity - information should be updated (values)	Workflow Efficiency	
Processes	Detection of Anomalous and Routine Events	Risk and Conjunction Assessment	Information Reception and Fusion	Sensor Allocation (Tasking Orders)	Maintenance of Asset State (Catalog Maintenance)	Degradation of Information Accuracy (Including TLEs) over Time
	UCT Processing	Release of Warnings and Information Dissemination				
Components	Individual Asset Catalog Records	Computational Resources (Supercomputers)	Total Space Asset Catalog	Sensor Systems	Spacetrack.org	Sensor Data (Raw)
Attributes	Number of objects being tracked and maintained in the catalog	Number of high-priority assets being tracked (1,300)	Number of radars	Number of electro-optical sensors	Measurement Uncertainty	

A work-domain model consists of a boundary that provides a scope for the socio-technical system being studied. The boundary for our analysis can be circumscribed around the sensor network and resources, earths atmosphere, and regions in the space environment that are pertinent to the safety of monitored space assets. This may require two abstraction hierarchies, in addition to which a third abstraction hierarchy for the space surveillance network would be useful as a separate entity with closely shared. Such a work-domain model consisting of multiple, related abstraction hierarchies has been reported by Burns, Bryant, and Chalmers (2005).

In this paper, we report a draft abstraction hierarchy for a SSA decision-support environment, as shown in Table 1. The work-domain model presented here is based on interviews with subject-matter experts with extensive experience in SSA decision-support environments and processes involved in maintaining SSA. The highest level of the abstraction hierarchy in Table 1 describes the purposes of SSA decision-support systems, which include space asset safety and security, national security, and workflow efficiency. These purposes are achieved through effective prioritization of critical assets, expectations generated about space objects through the knowledge of physics, and a sufficient understanding of uncertainty in the environment. Several processes are in place to support expectations and the minimization of uncertainty. These processes can be maintained by various resources listed in the fourth row referred to as components, which have specific attributes listed in the lowest level.

2.2. Hypothesis Abstraction

Many existing sensor tasking approaches perform well in maintaining a low overall uncertainty (e.g. information-maximum), but this tasking does not necessarily support the needs of a decision-maker. It may not be readily apparent to a decision-maker or




Figure 1: Predictive and reactive sensor tasking.

sensor-tasker how reducing a specific object’s uncertainty by a certain amount affects situation awareness or answers decision-making questions. This motivates an approach that encodes tasking opportunities and decision-making priorities as hypotheses that can be interrogated by evidence. If a potential target’s orbit and operational capabilities (or lack thereof) are well-known, it might not be necessary to minimize its associated uncertainty. Conversely, many consecutive follow-ups might be desired on a newly-acquired object to fully characterize its orbit, or on an object approaching a congested volume of space (such as a GTO object approaching apogee).

Hypothesis-driven approaches are not new to SSA; for instance, multiple hypothesis testing (MHT) techniques have been applied to object detection within electro-optical images [32, 33, 34]. Applied to sensor-tasking, hypothesis-driven approaches enable a predictive mode of tasking to answer specific relevant questions. This is fundamentally different from reactive sensor tasking approaches, where the gathered information is used to form hypotheses about what caused the observed behavior. Predictive sensor tasking uses prior knowledge of relevant hypotheses to estimate information-gain from potential courses of action. This relationships between hypotheses and information in predictive and reactive sensor tasking are illustrated in Fig. 1.

Re-framing sensor tasking in terms of hypotheses supports human cognition in decision-making. While automation excels at performing computation-heavy analyses such as data reduction and processing, humans excel at more abstract-level cognitive tasks required for objective prioritization and goal-adjustment [9]. This corresponds to the top-levels of the abstraction hierarchy in Table 1, while successively lower levels refer to finer-detailed portions of the levels above. Forcing an operator to switch between different levels of the abstraction hierarchy, effectively approaching the problem at multiple different levels of detail, leads to increased frustration and workload and decreased situation awareness. Designing a decision-support system that directly conveys hypothesis resolution information ensures that the human decision-maker spends more time on strategic cognitive tasks.

2.3. Ambiguity Aversion

Multiple methodologies exist for modeling and reasoning in uncertain domains to provide graphical and numerical representations of uncertainty [35]. One prevailing methodology is Bayesian probability theory, which models knowledge about propositions using true-or-false probabilities. However, probability theory struggles to express ambiguity in proposition knowledge, often due to some ignorance on the part of the expert or evidence source.

For illustration, consider an expert with vacuous knowledge, or total ignorance, on a proposition. In probability theory, this is often represented using the principle of non-information: each state in the proposition state space is assigned equal probability. This equally-likely probability mass function can also arise naturally when an expert

has certain knowledge that places equal probability on each state. Therefore, the same probability mass function can represent two very different knowledge states, one with wholly ambiguous information and the other with certain but conflicting evidence, due to the inability to encode ambiguity in Bayesian probability [35]

It has been shown that human decision-makers overwhelmingly prefer known risks to unknown risks, making ambiguity a major concern in modeling knowledge states [36]. Ellsberg’s paradox, re-stated here, is a well-known example that violates Savage’s theory of subjective expected utility [36]. Consider two urns, each filled with 100 red or yellow balls. The first urn contains an unknown distribution of red and yellow balls. The second urn contains an equal distribution of red and yellow balls, 50 of each. The goal is to draw a red ball from one of the urns, and the human decision-maker is allowed to choose which urn they draw from. The results of Ellsberg’s study show that humans overwhelmingly chose to draw from the second urn, which has a known probability distribution, even though the first urn may contain a favorable distribution of red balls. This is a phenomenon known as “ambiguity aversion” and is a predictable characteristic of human decision-making in the face of uncertainty.

The first urn in Ellsberg’s paradox represents a vacuous knowledge state, while the second urn represents the equal-probability knowledge state. Using Bayesian probability, both knowledge states would be represented with the same probability mass function, meaning the information presented to the decision-maker would not adequately convey information on the presence or lack of ambiguity that would impact the decision. This highlights a deficiency in Bayesian probability theory that has a significant impact in human decision-making contexts, which motivates the use of alternative methodologies such as evidential reasoning. One of the most prevalent alternatives to Bayesian probability is Dempster-Shafer theory, and the relevant aspects of the theory are presented in the following section.

2.3.1. Dempster-Shafer Theory Background

Dempster-Shafer theory is considered more expressive than probability theory in representing ambiguity or ignorance [37]. This is accomplished by allowing assignment of belief to non-singleton propositions, admitting ambiguity on the part of the expert when necessary.

In Dempster-Shafer theory, the possible propositions of a given hypothesis form a set called the frame of discernment, Ω . The frame must be a set of mutually exclusive and collectively exhaustive propositions [19], though some alternative formulations such as the Transferable Belief Model allow for relaxation of the latter constraint [38]. A basic belief assignment (BBA), as defined in Eqn. (1), maps a belief mass to each possible proposition:

$$m : \theta \mapsto [0, 1], \quad \theta \in 2^\Omega \quad (1)$$

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

$$m(\emptyset) = 0 \quad (3)$$

where m is the BBA and 2^Ω is the set of all subsets of Ω . Elements of 2^Ω are referred to as propositions. Note that, notationally, $\{\theta_1, \theta_2\}$ is equivalent to $\theta_1 \cup \theta_2$. The constraint in Eqn. (2) enforces the mutually exclusive and collectively exhaustive properties as the belief masses must sum to one, while the constraint in Eqn. (3) is similar to Kol-

mogorov's axiom of zero probability for the empty set. Allowing attribution of belief mass to non-singleton propositions enables encoding of evidence ambiguity.

Using BBAs, Shafer defines notions of belief (or support) and plausibility, which form lower and upper bounds respectively on the probability that a proposition is true given the available evidence [39]. The belief in, or support for, proposition $A \in 2^\Omega$ is defined in Eqn. (4) as the sum of belief masses attributed to A and its proper subsets.

$$\text{Bel}_m(A) = \sum_{\theta|\theta \subseteq A} m(\theta) \quad (4)$$

The plausibility of proposition $A \in 2^\Omega$ is defined in Eqn. (5) as the sum of the belief masses that do not explicitly refute proposition A :

$$\text{Pl}_m(A) = \sum_{\theta|\theta \cap A \neq \emptyset} m(\theta) = 1 - \text{Bel}_m(\neg A) \quad (5)$$

where $\neg A = 2^\Omega \setminus A$ is the negation of hypothesis A , or "not A ." The BBA representations of belief mass, belief, and plausibility are all interchangeable via the above linear relationships [40].

Various BBA combination rules have been developed to fuse evidence from multiple sources into one BBA [41]. The most common combination rule is Dempster's conjunctive rule [39], defined in Eqn. (6):

$$m_{1 \oplus 2}(A) = (m_1 \oplus m_2)(A) = \sum_{B, C \subseteq \Omega | B \cap C = A} (1 - K)^{-1} m_1(B) m_2(C) \quad (6)$$

$$K = \sum_{B, C \subseteq \Omega | B \cap C = \emptyset} m_1(B) m_2(C) \quad (7)$$

where K , defined in Eqn. (7) is a normalization conflict that quantifies internal conflict in the BBAs.

2.3.2. Decision-Making with Dempster-Shafer Theory

While the ability to represent ambiguity in belief functions is useful for accurately representing knowledge states, the theory of belief functions lacks a coherent decision theory [35]. Multiple methods exist for translating between Dempster-Shafer belief functions and probability models, allowing the use of Bayesian decision theory. Smets suggested the use of the pignistic transformation [38], but it has been argued that the pignistic transformation may not be consistent with Dempster's rule of combination [35]. An alternative method, the plausibility transformation, is defined in Eqn (8) [35]:

$$\text{Pl}_{P_m}(x) = K^{-1} \text{Pl}_m(\{x\}) \quad (8)$$

$$K = \sum_{x \in \Omega} \text{Pl}_m(\{x\}) \quad (9)$$

Note that the normalization constant K in (9) is different from the normalization constant for Dempster's conjunctive rule in Eqn. (7). The plausibility transformation is consistent with Dempster's rule, particularly in situations where pignistic probability is inconsistent [35].

Another important concept in both probabilistic and evidential reasoning is entropy as an information content measure. For Dempster-Shafer theory, multiple definitions of

Table 2: Ellsberg’s paradox belief structures and entropy

Urn	$m(\{\text{red}\})$	$m(\{\text{yellow}\})$	$m(\{\text{red, yellow}\})$	$H_S(m)$	H_{DP}	H_{JS}
1	0.5	0.5	0	1	0	1
2	0	0	1	1	1	2

entropy have been proposed, many of which are summarized by Jirousek and Shenoy [37]. Conflict in the belief structure is measured through Shannon entropy using the plausibility transform, where low conflict means a significant belief mass attributed to a singleton proposition. Non-specificity captures ambiguity as the entropy associated with non-singleton focal sets of the bba using the Dubois-Prade entropy. The Jirousek-Shenoy (J-S) definition of entropy combines Shannon and Dubois-Prade entropy to capture both conflict and non-specificity. Minimizing both conflict and non-specificity ensures that the resulting belief structure is internally consistent (i.e. prefers strong hypothesis resolution over an equally-probable result) and is non-ambiguous.

One useful property of J-S entropy is that maximum entropy is only attained by a vacuous bba, which is the bba where all belief mass is assigned to the entire frame: $m(\Omega) = 1$. Including both conflict and non-specificity (or ambiguity) in the entropy calculation allows for appropriate modeling of the ambiguity aversion phenomenon [37]. Recalling Ellsberg’s paradox, the first urn is an equally-likely belief structure and the second urn is a vacuous belief structure:

$$\begin{aligned} m_1(\{\text{red}\}) &= m_1(\{\text{yellow}\}) = 0.5, & m_1(\{\text{red, yellow}\}) &= 0 \\ m_2(\{\text{red}\}) &= m_2(\{\text{yellow}\}) = 0, & m_2(\{\text{red, yellow}\}) &= 1 \end{aligned}$$

The Shannon entropy, Dubois-Prade entropy, and J-S entropy for these belief structures are shown in Table 2. As expected, Shannon entropy shows high conflict for both belief structures, but Dubois-Prade entropy is only non-zero for the ambiguous distribution, so the second urn has a higher J-S entropy. The decision-maker wants to minimize conflict and non-specificity, so selecting urn 1 with the lower J-S entropy is consistent with the result from Ellsberg’s paradox. Therefore, minimizing J-S entropy can be used as a reliable and consistent metric for a strong hypothesis resolution.

2.4. Unbiased Hypothesis Resolution

Hypothesis resolution refers to the goal of determining which proposition is true from the set of propositions in the frame of discernment. The previous section discussed the use of Dempster-Shafer belief structures to determine when a strong resolution is reached. However, one phenomenon worth considering in any hypothesis resolution task is confirmation bias. Confirmation bias is a human cognitive phenomenon where a prior belief causes fixation on a particular proposition, causing the human to favor evidence that confirms prior beliefs and overlook conflicting evidence. In regimes where uncertainty and ambiguity are a concern, this effect also applies by interpreting ambiguous evidence in favor of the prior beliefs.

A non-human system such as a sensor network might also exhibit confirmation bias as prior information has the potential to skew future evidence-gathering actions. A most-probable-first sensor tasking approach would prioritize actions that gather further evidence to confirm the prior. Similar to human cognitive fixation, a technological fixation may result from incorrect prior assumptions or ambiguous evidence induced by

measurement noise, sensor bias, or other sources of uncertainty. For instance, spurious detections or false alarms may lead increased belief in the incorrect proposition. For illustration, consider a binary frame $\Omega = \{x_1, x_2\}$ and a prior that places slight belief in the x_1 proposition: $m(\{x_1\}) = 0.1$, $m(\{x_2\}) = 0$, $m(\{x_1, x_2\}) = 0.9$. A most-probable-first sensor tasking algorithm would then focus further actions on confirming the incorrect proposition, which ignores the (much larger) ignorance in the estimated proposition.

It is important to avoid fixating on any particular proposition where incorrect priors or evidence ambiguity may be the cause of any bias, adding a competing objective to the requirement of minimizing hypothesis entropy. An algorithm that is only concerned with resolving hypotheses in minimum time would prefer the most-probable tasking approach and may suffer from confirmation bias. However, development of a reliable method for detecting and avoiding confirmation bias would be highly problem dependent: such a metric depends on specifics including the types of evidence and evidence sources, known or calibrated noise and bias characteristics, how the evidence maps to the hypotheses, and potentially many other problem specific parameters. Just as fixation should not be ignored in favor of time optimality, fixation should not be the only focus at the cost of resolving hypotheses within time constraints. Rigorous and time-efficient techniques for quantifying confirmation bias from a proposed tasking schedule are likely not to be generalized, so instead an alternating-turns heuristic is applied to allocate equal time to confirming and refuting each hypothesis.

An apt human decision-making analogy for this heuristic is the fair trial system, wherein opposing agents (defense and prosecution) are given equal opportunity to present the strongest evidence to confirm or refute a hypothesis. Each agents must select evidence to prove its assigned proposition, and at the end of the decision horizon (the trial), the sum-total of the evidence is fused to render a decision based on the proposition with the most belief. The fair trial system is by no means time optimal, but the presence of competing agents ensures that no one side is allowed to dominate the narrative of the case. This encourages appropriate resolution as guided by the evidence, not prior beliefs, biases, or ambiguity.

Similarly, the proposed unbiased hypothesis resolution framework employs an agent for each proposition of the hypothesis, and these agents alternate turns to allow equal opportunity for gathering supporting evidence. Application of this alternating-turns heuristic is a result of the trade between time-optimality and unbiased hypothesis resolution, and due to strong parallels to the fair trial system, the proposed hypothesis resolution technique developed in this work is called Judicial Evidential Reasoning (JER).

2.5. Judicial Evidential Reasoning

The three primary considerations of the JER framework, as described above, are: hypothesis abstraction, ambiguity aversion, and unbiased hypothesis resolution. Employing a hypothesis abstraction enables predictive tasking and supports human cognition at a strategic and planning level. The use of evidential reasoning, specifically Dempster-Shafer theory, to model hypothesis knowledge allows quantified ambiguity and ambiguity together in the entropy measurement. Applying the alternating-turns heuristic, inspired by the fair trial system, provides ensures impartial or unbiased hypothesis resolution to guard against confirmation balance while still balancing a need for time-efficient hypothesis resolution. This section further develops the JER framework into an algorithmic implementation.

This formulation uses receding horizon optimization, targeting time t_{k+H} which is H steps after the current time t_k . Begin by defining the sets of considered hypotheses as follows:

$$\mathbb{H}_i : \text{hypotheses at time } t_i \quad (10)$$

$$\mathcal{H}_j \in \mathbb{H}_i \forall j \in 1, \dots, |\mathbb{H}_i| \quad (11)$$

$$\mathbb{W}_i : \text{weights corresponding to hypotheses } \mathbb{H}_i \text{ at time } t_i \quad (12)$$

$$w_j \in \mathbb{W}_i \forall j \in 1, \dots, |\mathbb{W}_i|, \quad \sum_{j \in 1, \dots, |\mathbb{W}_i|} w_j = 1 \quad (13)$$

where the weights are assigned by decision-makers according to priority of the related hypotheses. Similarly, sets of actions and action sequences, or actions taken over a time-span, by all sensors in the sensor network must also be defined. The available and selected sets of actions are defined as follows:

$$\mathbb{A}_i : \text{actions available at time } t_{k+i} \forall i \in 1, \dots, H \quad (14)$$

$$\mathcal{A}_i : \text{actions selected at time } t_{k+i} \forall i \in 1, \dots, H \quad (15)$$

$$\mathcal{A}_i \subseteq \mathbb{A}_i \quad (16)$$

The available and selected sequences of actions over the time horizon are defined as follows:

$$\mathbb{A}_{1,H} : \text{action sequences available over the time horizon } [t_{k+1}, t_{k+H}] \quad (17)$$

$$\mathbb{A}_{1,H} = \mathbb{A}_1 \times \dots \times \mathbb{A}_H \quad (18)$$

$$\mathcal{A}_{1,H} : \text{action sequences selected over the time horizon } [t_{k+1}, t_{k+H}] \quad (19)$$

$$\mathcal{A}_{1,H} = \mathcal{A}_1 \times \dots \times \mathcal{A}_H \quad (20)$$

$$\mathcal{A}_{1,H} \in \mathbb{A}_{1,H} \quad (21)$$

The multiple-hypothesis optimization problem aims to determine the action sequence ($\mathcal{A}_{1,H}^*$) to arrive at an unbiased resolution for all the hypotheses, prioritized by the hypothesis weights. This is represented generically by minimizing the following objective function at time t_{k+H} :

$$\mathcal{A}_{1,H}^* = \arg \min_{\mathcal{A}_{1,H} \in \mathbb{A}_{1,H}} J_H(\mathbb{H}_H, \mathbb{W}_H; \mathcal{A}_{1,H}) \quad (22)$$

$$J_H : (\mathbb{H}_H, \mathbb{W}_H; \mathcal{A}_{1,H}) \mapsto \mathbb{R} \quad (23)$$

Due to the combinatorial nature of this problem, evaluating all possible action sequences over a given time horizon exhibits exponential complexity growth. For this reason, brute-force solutions quickly become intractable and a method for efficiently selecting action subsequences is preferred. Recalling the desired alternating-turns heuristic, the multi-hypothesis problem is decomposed into a sub-problem for each hypothesis.

2.5.1. JER Hypothesis Agents

Consider a single hypothesis from the set of considered hypotheses at time t_k : $\mathcal{H} \in \mathbb{H}_k$, where $\Omega_{\mathcal{H}}$ is its frame of discernment and contains $|\Omega_{\mathcal{H}}|$ propositions. Each

proposition must be either conclusively confirmed or refuted with evidence, so each proposition is assigned a pair of JER agents:

$$a_{\mathcal{H}}(\theta), a_{\mathcal{H}}(-\theta) \quad \forall \theta \in \Omega_{\mathcal{H}} \quad (24)$$

where $-\theta = \Omega_{\mathcal{H}} \setminus \theta$. Therefore, for hypothesis \mathcal{H} there are $|\Omega_{\mathcal{H}}|$ alternating JER agent-pairs. Each agent pair is initialized with the $(a_{\mathcal{H}}(\theta))$ agent as the active agent.

For each agent-pair, compute the optimal action sequence over the time horizon by evaluating possible action sequences. One possible computational tractability improvement is to limit the action sequences evaluated to only those which will impact this agent-pair's hypothesis. Since evaluating potential action sequences exhibits exponential growth in complexity with added actions or increased time horizon, reducing the set of possible actions evaluated may significantly improve computational complexity.

Since the goal for the null agent is to maximize belief in the null proposition θ , this can be formulated as a maximin optimization using the plausibility transformation:

$$\mathcal{A}_{1,H|a_{\mathcal{H}}(\theta)}^* = \arg \max_{\mathcal{A}_{1,H} \in \mathbb{A}_{1,H}} \min \text{Pl}_P(\theta; m_{\mathcal{H}|\mathcal{A}_{1,H}}) - 0.5 \quad (25)$$

where $m_{\mathcal{H}|\mathcal{A}_{1,H}}$ is the estimated BBA resulting from the proposed action sequence $\mathcal{A}_{1,H}$. The plausibility transformation is applied here because of its relationship and consistency with decision-making. Additionally, the quantities $\text{Pl}_P(\theta)$ and $\text{Pl}_P(-\theta)$ always sum to one, so to make the formulation a zero-sum game (as per convention in adversarial optimization), the 0.5 term is subtracted from the maximin objective. The maximum attainable value for this objective is 0.5 when proposition θ has full belief, and the minimum attainable value for this objective is -0.5 when proposition $-\theta$ has full belief.

When the alternative agent $(-\theta)$ is active, its goal is to maximize belief in the alternative proposition or equivalently minimize belief in the null proposition. Therefore, the formulation simply flips to a minimax optimization:

$$\mathcal{A}_{1,H|a_{\mathcal{H}}(-\theta)}^* = \arg \min_{\mathcal{A}_{1,H} \in \mathbb{A}_{1,H}} \max \text{Pl}_P(\theta; m_{\mathcal{H}|\mathcal{A}_{1,H}}) - 0.5 \quad (26)$$

The result of the JER agent-pair schedule optimization is an optimal action sequence for the active agent. If that agent-pair's action is selected, that agent-pair flips active agents for the next time step.

2.5.2. Resolving Combined Schedule Incongruity

After optimal schedules for each JER agent-pair from each hypothesis sub-problem are evaluated, the schedules must be combined into a single schedule. Depending on the hypotheses, it is possible or even likely that two different agent-pairs will compute optimal sequences that use the same sensor for different actions, representing an incongruity between the schedules. These are resolved by choosing the actions that lead to the strongest hypothesis resolution using entropy.

Using the set of actions from all sub-problem optimal sequences $\mathcal{A}_{1,H|a_{\mathcal{H}}(\theta)}^*$ all possible combinations of these actions are used to propose congruous action sequences. At worst case, this is the same as a brute-force re-evaluation, but that would require that all hypotheses have the same applicable action subsets and all possible actions produce an optimal result for at least one hypothesis. This implies an extreme interdependence between the hypotheses that is unlikely to occur in operation. In more

realistic cases, where most hypotheses are distinct enough to have different applicable actions, this re-evaluation is much less computationally complex than brute-force.

The evaluation criterion for resolving incongruity between sub-problem action sequences is the weighted-sum of J-S entropy.

$$\mathcal{A}_{1,H}^* = \arg \min_{\mathcal{A}_{1,H}^* \in \mathcal{H}^{(0)}} \sum_{i=1}^{|\mathbb{H}_H|} w_i H_{JS} (m_{\mathcal{H}_i | \mathcal{A}_{1,H}}) \quad (27)$$

where H_{JS} is the J-S entropy. Since J-S entropy quantifies both conflict and non-specificity, and the weighting parameters encode decision-maker priorities, the resulting action sequence $\mathcal{A}_{1,H}^*$ is the action sequence with the strongest priority-weighted resolution.

2.5.3. Stopping Criterion

Once sufficient evidence has been gathered to resolve a hypothesis, it would be beneficial to remove that hypothesis from consideration for future tasking evaluations. Pruning the considered hypotheses ensures that sensor resources are not used to gather further evidence on an already-resolved hypothesis, and also reduces computational complexity because fewer JER agent-pair sub-problems need to be evaluated. As discussed in previous sections, a hypothesis is considered resolved when there is low conflict and low ambiguity. Therefore, an appropriate stopping criterion is the J-S entropy.

A J-S entropy threshold can be developed based on desired levels of hypothesis resolution, and can be customized for each hypothesis. For instance, decision-makers may want to ensure that a conjunction assessment is correct before sending conjunction and maneuver notices, necessitating tighter thresholds for conjunction hypothesis resolution than other hypotheses. Given a desired maximum amount of conflict as characterized by Shannon entropy $H_{S,max}$, and a maximum level of non-specificity as characterized by Dubois-Prade entropy $H_{DP,max}$, the threshold value of Jirousek entropy is $H_{JS,max} = H_{S,max} + H_{DP,max}$. The specific values for these thresholds can be determined by constructing a threshold bba for each hypothesis $m_{\mathcal{H},th}$, which represents worst-case hypothesis knowledge for that hypothesis to be considered resolved, and computing the Shannon and Dubois-Prade entropies for that threshold bba. See Jirousek and Shenoy for the specific equations on computing each entropy [37].

2.5.4. JER Algorithm Summary

The JER framework developed in the previous sections is summarized in Fig. 2 to illustrate the complete algorithm flow.

3. Simulation Results

This section contains an illustrative example application of JER autonomously scheduling actions to resolve multiple hypotheses. The JER framework described above is implemented in Python for testing. Space objects are propagated using Keplerian two-body dynamics to compute lines-of-sight to sensors, illumination conditions using a cannonball model. Two electro-optical sensors are simulated, separated by 20 degrees in longitude for geometric diversity. Both sensors are 3-degree field-of-view search telescopes. Electro-optical observations are simulated using a radiometric model, including simulated effects for background sky irradiance and atmospheric transmittance




Figure 2: Judicial Evidential Reasoning algorithm

(e.g. cloud cover, atmospheric seeing conditions) [42]. The sensors may change actions each minute, and a receding time-horizon of three minutes is used. Hypotheses are considered resolved if the normalized Jirousek-Shenoy entropy drops below the threshold value of $H_{JS,thr} = 0.05$, requiring very low conflict and ambiguity.

A limited subset of potential failure modes is analyzed for illustrative purposes for this test case. Since multiple point-of-failure events are exceedingly rare, an assumption is made that the anomaly results from a single point-of-failure.

3.1. Scenario Description

Operators in a SSA decision-support environment receive notice from a space launch entity that a planned GTO insertion maneuver has experienced an anomaly. The anomaly is estimated to have occurred 5 minutes prior to the notification during a critical orbit-raising maneuver, after which no communications have been received. The objective is to re-acquire the space object and diagnose the anomaly to regain situation awareness.

Anomalous GTO objects are particularly difficult to characterize as the range prohibits use of radar, requiring a wide state-space search using electro-optical sensors. Timely re-acquisition is critical since the spacecraft was initially bound for Geostationary Earth Orbit (GEO), a densely populated orbit regime with many high-value defense and telecommunications assets. If the anomaly still resulted in a GEO-intersecting trajectory, it is crucial to characterize the new orbit to inform conjunction analyses. Similarly, if the resultant trajectory remains close to low-Earth orbit (LEO), it becomes a collision risk in a densely populated orbit regime.

The entire simulation occurs over a 15 minute time span, including the 5-minute delay between the anomaly event and the beginning of the sensor tasking window. The simulation time span is limited by observation constraints (e.g. short horizon-to-horizon

GTO Insertion Failure - Nominal Transfer




Figure 3: Nominal GTO transfer orbit and target GEO orbit.




Figure 4: Possible causes for GTO insertion failure

times in LEO, eclipse, adverse weather), placing time pressure on the hypothesis resolution. At the end of this simulation the sensor positions will prohibit gathering further evidence, so the anomaly must be characterized within 25 minutes of the event.

The nominal transfer orbit is shown in Fig. 3. The spacecraft begins in a 1000 km altitude circular parking orbit. The nominal transfer time from LEO to GEO is just over 5 hours, placing additional time-pressure on resolving the anomaly to complete conjunction analyses with time remaining to alert other satellite operators.

3.2. Belief Function Models

As shown in Fig. 4, the anomaly is characterized at the subsystem level to determine root-cause (under the stated single-point-of-failure assumption). The considered propositions for a GTO insertion maneuver include: maneuver status, propulsion status, navigation status, and collision in LEO. This decomposition yields the following top-level frame of discernment:

$$\Omega = \Omega_{man} \times \Omega_{prop} \times \Omega_{nav} \times \Omega_{coll} \tag{28}$$

To arrive at JER agent-pair frames, each element in frame Ω is further decomposed into sub-frames.

3.2.1. Maneuver Status

The maneuver status hypothesis, Ω_{man} , is simply decomposed into a binary frame for the nominal and anomalous states:

$$\Omega_{man} = \{\omega_{man,n}, \omega_{man,a}\} \quad (29)$$

Since it does not decompose further, Ω_{man} is assigned a single JER agent-pair to test whether the maneuver is nominal or anomalous. Evidence for the for the nominal maneuver status proposition includes a successful nominal detection of the space object in its intended geostationary transfer orbit. Additionally, evidence for the anomalous maneuver status proposition can be stated by an implicit relationship to the other subsystem frames: if any of these hypotheses is resolved true in an anomalous state, the anomalous maneuver status hypothesis is true.

3.2.2. Propulsion Status

The propulsion status hypothesis, Ω_{prop} , yields a more complex (non-binary) decomposition using known failure modes for the propulsion subsystem:

$$\Omega_{prop} = \{\omega_{prop,nom}, \omega_{prop,ns}, \omega_{prop,exp}\}$$

Nominal propulsion status, $\Omega_{prop,nom}$, is the proposition required to exhaust all possible propulsion subsystem states, since the propulsion subsystem might not be the cause of the anomaly. The non-start proposition, $\Omega_{prop,ns}$, occurs when the propulsion system fails to fire, leaving the spacecraft in its LEO parking orbit. The explosion proposition, $\Omega_{prop,exp}$, occurs when there is a catastrophic failure, resulting in debris in LEO near the spacecraft's parking orbit.

Evidence for the non-start proposition includes successful detection of the spacecraft in its LEO parking orbit, while evidence for the explosion proposition includes multiple detections of debris pieces along the LEO parking orbit. Note that debris could also be indicative of collision, so the explosion proposition must be refuted if debris is detected in multiple orbits.

3.2.3. Navigation Status

The navigation status hypothesis, Ω_{nav} , does not decompose into further complex sub-frames as the only proposition to consider is an off-nominal transfer orbit due to pointing error.

$$\Omega_{nav} = \{\omega_{nav,n}, \omega_{nav,a}\} \quad (30)$$

Since it does not decompose further, Ω_{nav} is assigned a single JER agent-pair to test whether the navigation status is nominal or anomalous.

3.2.4. Collision in LEO

The collision in LEO hypothesis, Ω_{coll} , decomposes into the following non-binary frame:

$$\Omega_{coll} = \{\omega_{coll,none}, \omega_{coll,1}, \dots, \omega_{coll,R}\}$$

where R is the number of resident space objects (RSOs) to consider for potential collisions. The "none" proposition, $\omega_{coll,none}$, represents the case where a collision has not

occurred and therefore is not the cause of the anomaly. Collision with object j , $\omega_{Coll,j}$ where $j = 1, \dots, R$, occurs when the spacecraft collides with object j in LEO, resulting in a debris in both orbits as well as missing nominal tracks for both object j and the primary spacecraft. Note that the conditional evidence mentioned in the propulsion status section handles this differentiation. Nominal detection of an RSO refutes that RSO's collision proposition.

For this illustrative example, three RSOs ($R = 3$) will be considered for close-approach and potential collision in LEO. Therefore, three JER agent-pairs are used to fully explore this hypothesis, one for each collision proposition.

3.2.5. JER Agent-Pairs

The full problem considers each frame described in the decomposition above to investigate the cause of a maneuver anomaly. Each frame contains $|\Omega_i| - 1$ JER agent-pairs: two for propulsion failure, one for navigation status, and three for collision in LEO. Therefore, there are a total of six JER agent-pairs in this simulation. The maneuver status hypothesis does not receive its own JER agent-pair because its evidence is modeled primarily as implicit evidence from the result of the subsystem analyses.

3.3. Evidence to Belief Function Mappings

Each potential action is evaluated for its estimated effect on the considered hypotheses to develop evidence-to-belief-function mappings. This process is highly problem-specific, requiring the modeler to consider what each potential successful or missed detection means with respect to each hypothesis. For instance, a missed detection of the nominal orbit may indicate anomaly, but if the estimated electro-optical probability of detection predicted a low chance of success, the belief mass should be attributed to ignorance instead.

Additionally, implicit knowledge about relationships between these frames can be imposed through conditional bbas [43]. In particular, it is known that, if evidence confirms that none of the subsystems are nominal, the maneuver status is likely nominal. A small chance (0.01) is allowed that there may be other causes for maneuver anomaly even if the modeled causes are nominal to account for mis-modeling of the problem. Similarly, if any one of the other causes is anomalous, then the maneuver status is likely anomalous. This is the implicit evidence used to relate the subsystem frames to the maneuver status frame.

3.4. Case 1: Nominal Maneuver

As a baseline, the true proposition for this case is the nominal maneuver status. The resulting sensor tasking schedule is shown in Fig. 5a, and Fig. 5b shows the normalized Jirousek-Shenoy entropy for each hypothesis.

At the first observation epoch, both sensors are only able to observe GTO, and since the maneuver was successful and observation conditions are good, it is detected right away. This immediately resolves both the propulsion status and navigation status hypotheses to "nominal" as shown in Figs. 6a and 6b, leaving only the collision hypotheses for investigation. Also note that collision object "LEO 2" passes close-enough to the GTO sky position to also be detected on this first step, resolving the relevant proposition for the collision in LEO hypothesis as shown in Fig. 6c.

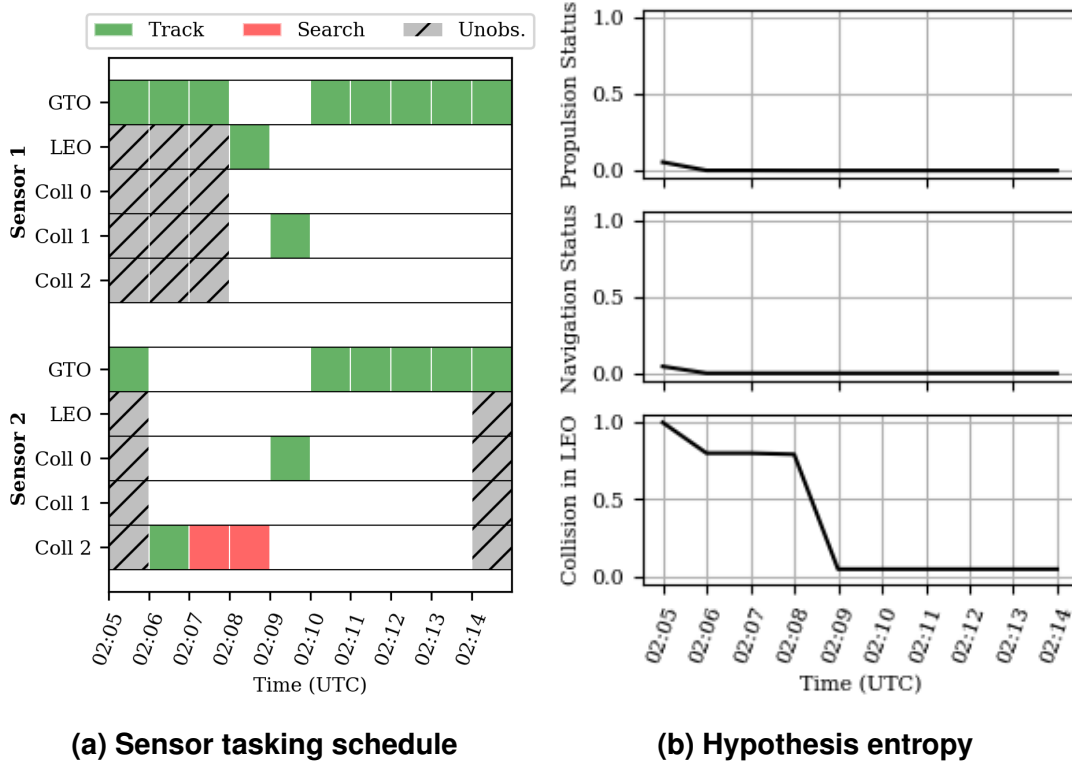


Figure 5: Case 1: nominal maneuver

As the sensors gain line-of-sight to the remaining collision objects, they are also observed, detected, and correlated. After five minutes of observations, all hypotheses have been resolved within the entropy threshold, so the remainder of the tasking actions are spent tracking the GTO object.

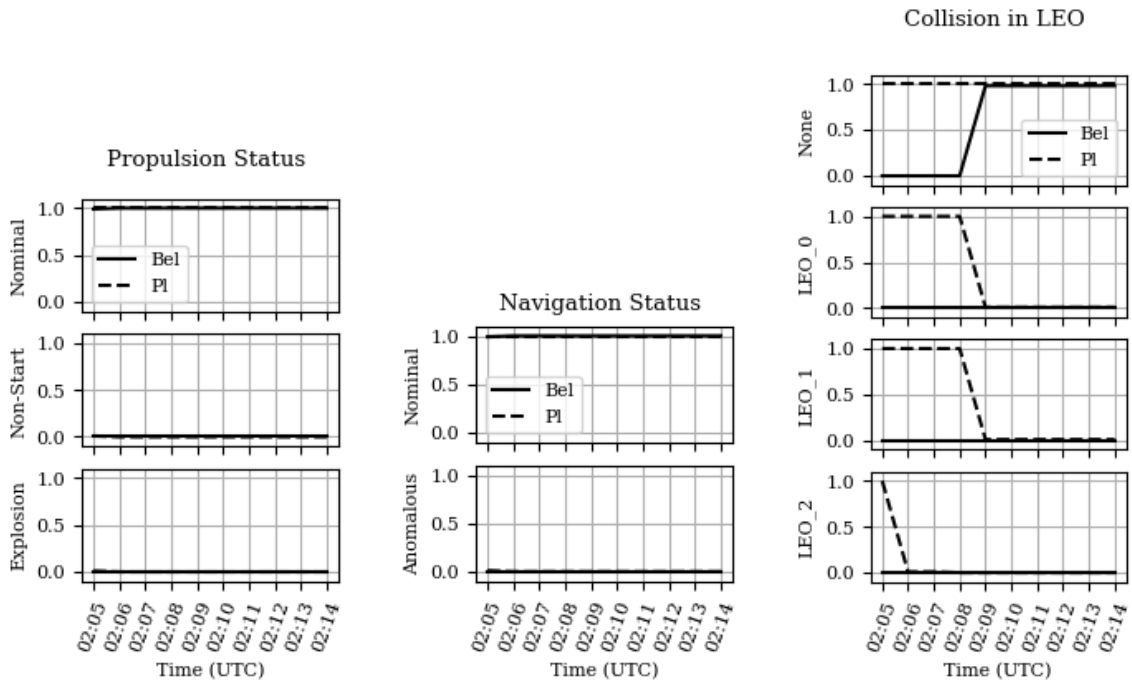
This baseline test case is intended as a comparison point of a nominal event. Even though the original notification of anomaly was incorrect (e.g. a false alarm), this does not bias the JER algorithm and the correct hypothesis resolution is obtained quickly. The following test cases will examine behavior in anomalous events.

3.5. Case 2: Navigation Anomaly

In the next test case, a navigation anomaly occurs resulting in a pointing error during the maneuver and a misaligned velocity vector. The resulting sensor tasking schedule and hypothesis entropies are shown in Figs. 7a, and 7b, respectively. Individual hypothesis propositions are plotted in Figs. 8a, 8b, and 8c.

Once again, the first observation epoch returns strong evidence for the correct hypothesis, but follow-up observations on the anomalous GTO object are required for several time-steps to ensure appropriate hypothesis resolution (recall the entropy threshold is conservative). This also serves the purpose of adequately resolving the propulsion status hypothesis to “nominal.” The final four observation epochs are used to observe the collision objects to refute collision in LEO. By the end of the simulation duration, all hypotheses have been resolved appropriately.

The JER algorithm narrows down on this proposition early since GTO is the only available tasking action at the beginning of the simulation, before the LEO objects are within line-of-sight of the sensors. The next test case involves a propulsion system anomaly and requires more time to discover the probable cause.




(a) Propulsion status

(b) Navigation status

(c) Collision in LEO

Figure 6: Case 1: nominal maneuver (baseline), hypothesis resolution



(a) Sensor tasking schedule

(b) Hypothesis entropy

Figure 7: Case 2: navigation anomaly

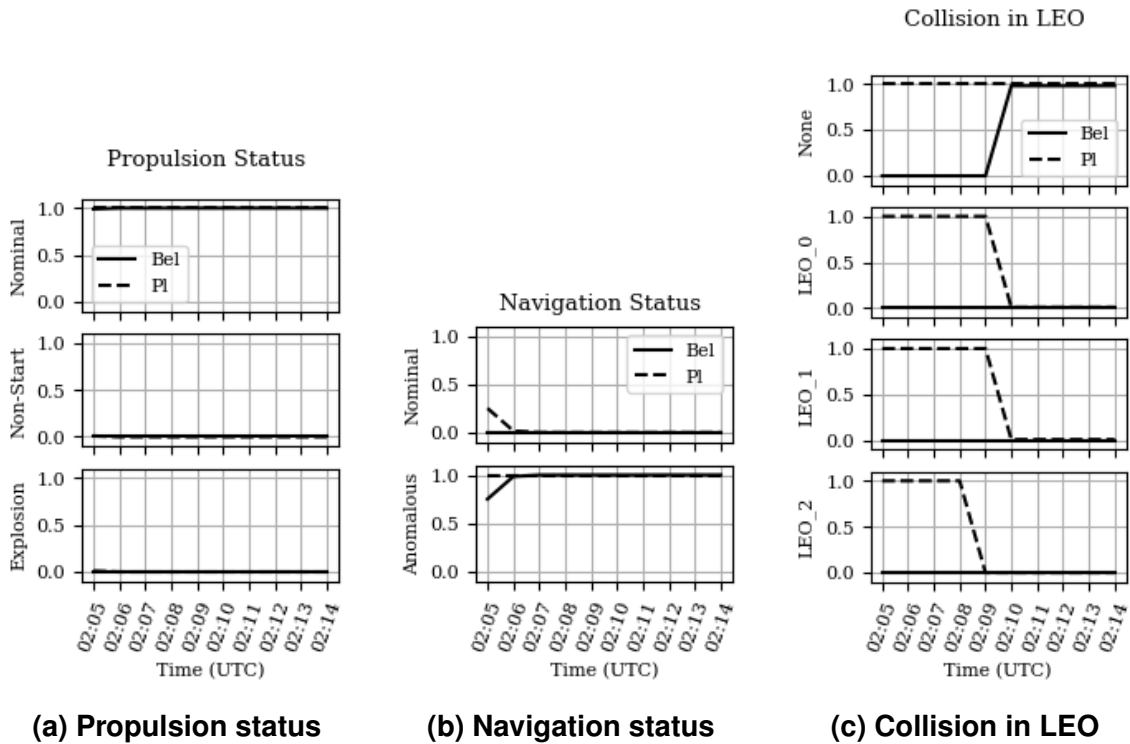


Figure 8: Case 2: navigation anomaly, hypothesis resolution

3.6. Case 3: Propulsion Non-Start

In the third test case, a propulsion anomaly occurs resulting in no maneuver and leaving the spacecraft in its LEO parking orbit. The resulting sensor tasking schedule and hypothesis entropies are shown in Figs. 9a, and 9b, respectively. Individual hypothesis propositions are plotted in Figs. 10a, 10b, and 10c.

At the beginning of this simulation, the missed detections at GTO only confirm the navigation status to be nominal after multiple actions allow exhaustive search of the near-GTO region. Once this hypothesis is resolved, the next actions quickly detect and correlate the spacecraft in LEO, as well as the collision objects. Since no debris is detected either, all hypotheses are resolved by the end of the fifth minute of simulation. No further tasking is required beyond this point.

This test case shows an ability to ingest both weak evidence (missed detections from GTO) and strong evidence (successful detections in LEO) to explore the hypotheses. The final test case presents a more challenging debris-generating event.

3.7. Case 4: Collision with Object in LEO

For the final test case, the true proposition is a collision in LEO with the object labeled “LEO 1.” This event generates multiple debris objects in both the LEO parking orbit and the nominal orbit of the collision object. The resulting sensor tasking schedule and hypothesis entropies are shown in Figs. 11a, and 11b, respectively. Individual hypothesis propositions are plotted in Figs. 12a, 12b, and 12c.

Once again, the GTO missed-detections quickly resolve the navigation status hypothesis. Next, observations of the collision object orbits only detect and correlate two of the objects, so belief in the collision proposition for LEO 1 begins to increase. Subsequent observations detect multiple debris objects in both the LEO parking orbit and

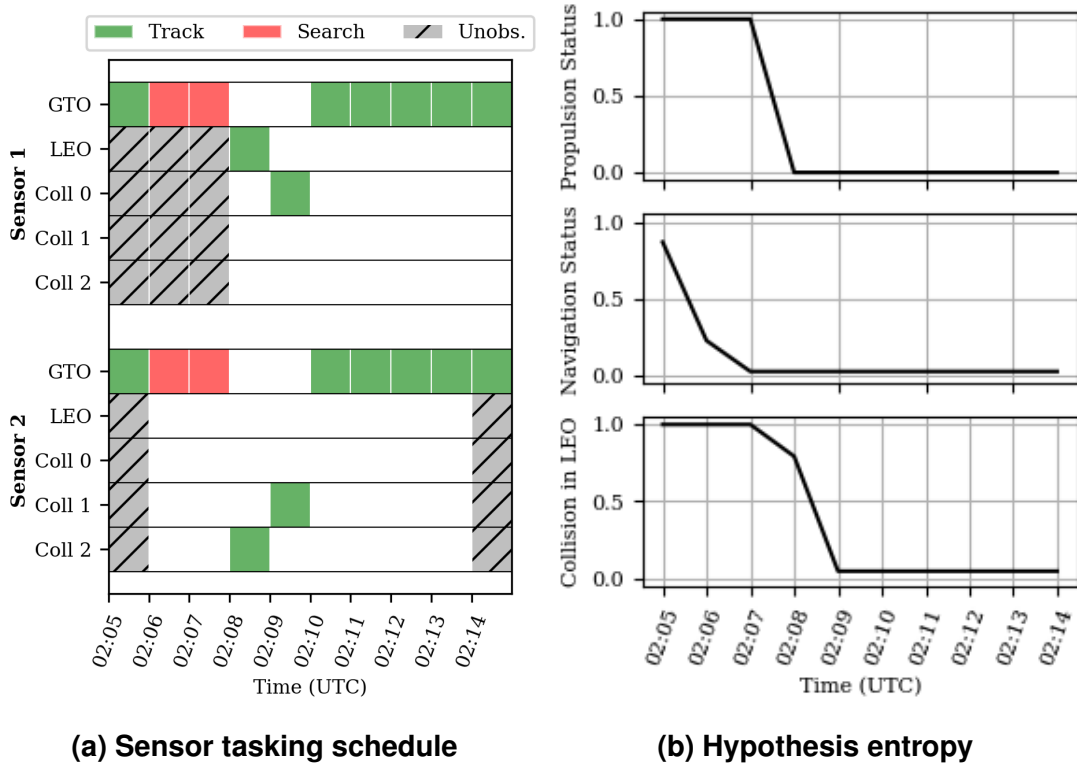


Figure 9: Case 3: propulsion non-start

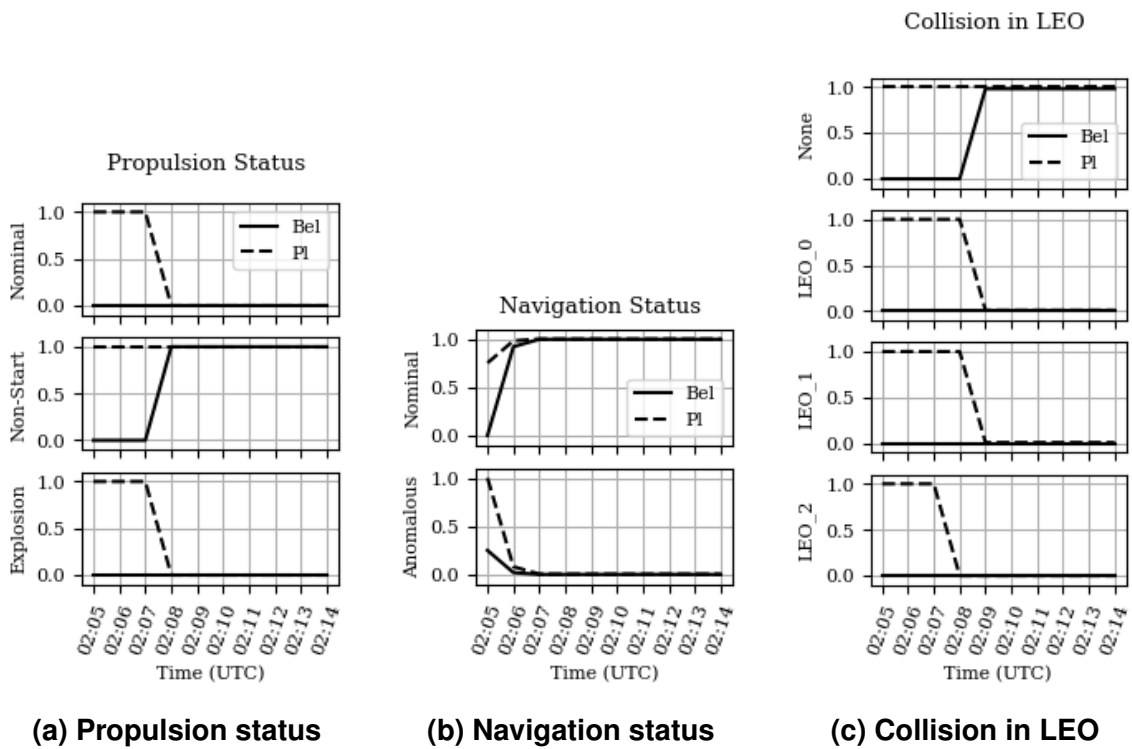


Figure 10: Case 3: propulsion non-start, hypothesis resolution

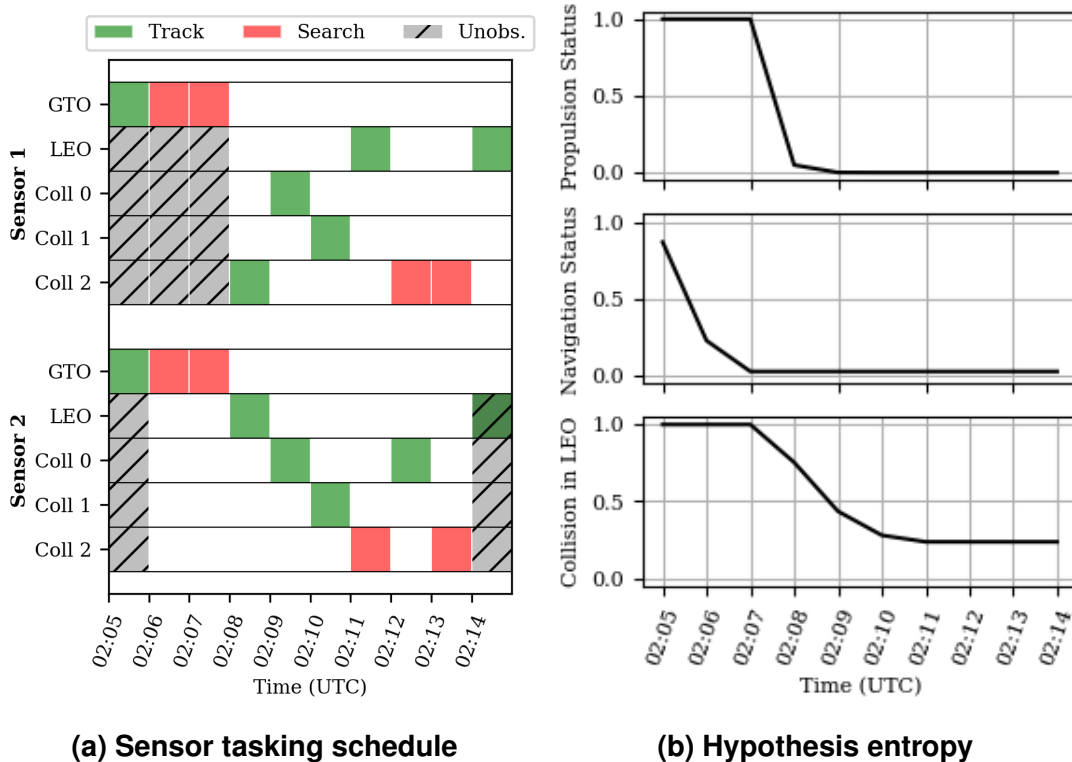


Figure 11: Case 4: collision in LEO

the nominal orbit of collision object LEO 1, while continued missed-detections of the nominal orbit for LEO also indicate the collision proposition. Importantly, the debris in the LEO parking orbit are not mistaken for explosion evidence due to the conditional evidence that links these two hypotheses.

At the end of this simulated case, two of the three hypotheses have been fully resolved, and the third (collision in LEO) has strong indication of resolution to the correct proposition. This last hypothesis is not fully resolved within the entropy tolerance because of conservative belief values applied to the detection of debris. Each new debris detected yields relatively low belief mass (0.25) in the relevant collision or explosion hypotheses, since real-world application would involve many spurious detections and false alarms caused by other resident space objects. Given more time or more sensors, the algorithm would detect more debris in both orbits, leading to increased belief and reduced entropy in the collision proposition.

3.8. Discussion

These simulated cases show that the JER algorithm performs as designed, seeking strong evidence to resolve hypotheses without fixating on any particular proposition. Weak evidence from missed detections results in the algorithm moving to other hypotheses or propositions that will plausibly produce stronger evidence. Additionally, decomposing the sensor tasking problem into tractable sub-problems through JER agent-pairs increases the feasible time horizon, which is computationally constrained in a brute-force approach, even for this relatively low-dimensional example.




Figure 12: Case 4: collision in LEO, hypothesis resolution

4. Conclusion

The proposed Judicial Evidential Reasoning (JER) sensor-tasking framework arranges decision-maker questions as rigorously testable hypotheses and employs an alternating-agent minimax optimization on belief in the null proposition. The use of a hypothesis abstraction supports human decision-making strengths of planning and strategy, off-loading processing work to the algorithm and fusing evidence into intuitive hypothesis resolutions. Recognizing the need to account for ambiguity aversion in decision-making, the use of Dempster-Shafer theory allows for quantification of evidence ambiguity. Finally, applying an alternating-turn adversarial optimization scheme avoids confirmation bias induced by improper prior beliefs or evidence uncertainty and ambiguity, avoiding fixation on incorrect propositions.

This approach values impartiality in addition to time-efficiency in many-hypothesis resolution, while breaking the greater sensor-tasking problem into a number of sub-problems for each hypothesis reduces computational complexity and allows for a receding horizon optimization of the total schedule. Selecting the final optimal schedules as the minimum total weighted entropy ensures that the selected actions minimize conflict and non-specificity according to priorities set by the decision-makers. The simulated results for a GTO insertion maneuver anomaly scenario show that the algorithm performs as expected: the appropriate hypotheses are confirmed via evidence and in the process the JER algorithm does not fixate on any particular proposition, instead accruing evidence that gradually leads to the correct conclusion.

Continuing work focuses on applying combinatorial and adversarial optimization approaches to the JER algorithm to prune sub-optimal sequences before deep evaluation, allowing application to larger multi-hypothesis problems with longer time horizons. Additionally, the need to include conditional evidence that explicitly links hy-

potheses (such as the explosion and collision hypotheses) motivates an approach using marginal belief functions [43].

References

- [1] M. R. Endsley, Situation awareness global assessment technique (SAGAT), in: IEEE 1988 National Aerospace and Electronics Conference.
- [2] J.-C. Liou, Modeling the Large and Small Orbital Debris Populations for Environment Remediation, Technical Report, NASA Orbital Debris Program Office, Johnson Space Center, Houston, TX, 2014.
- [3] Orbital Debris Quarterly News, Technical Report, NASA Orbital Debris Program Office, 2017.
- [4] M. Hart, M. Jah, D. Gaylor, B. C. T. Eyck, E. Butcher, E. L. Corral, R. Furfaro, E. H. Lyons, N. Merchant, M. Surdeanu, R. L. Walls, B. Weiner, A new approach to Space Domain Awareness at the University of Arizona, in: NATO Symposium on Considerations for Space and Space-Enabled Capabilities in NATO Coalition Operations.
- [5] J. D. Ianni, D. L. Aleva, S. A. Ellis, Overview of Human-Centric Space Situational Awareness Science and Technology, Technical Report, Air Force Research Lab, Human Performance Wing (711th), Human Effectiveness Directorate, Wright-Patterson Air Force Base, OH, 2012.
- [6] J. D. Rendleman, S. M. Mountin, Responsible SSA cooperation to mitigate on-orbit space debris risks, in: 7th International Conference on Recent Advances in Space Technologies (RAST).
- [7] T. Kelso, Analysis of the Iridium 33 Cosmos 2251 collision, in: Advanced Maui Optical and Space Surveillance Technologies, Maui, HI.
- [8] P. A. Brown, Promoting the safe and responsible use of space: Toward a 21st century transparency framework, *High Frontier 5* (2008).
- [9] F. R. Hoots, P. W. Schumacher, History of analytical orbit modeling in the U.S. space surveillance system, *Journal of Guidance, Control, and Dynamics* 27 (2004) 174–185.
- [10] T. Blake, M. Sanchez, J. Krassner, M. Georgen, S. Sundbeck, Space Domain Awareness, Technical Report, DARPA, 2012.
- [11] G. A. McIntyre, K. J. Hintz, An information theoretic approach to sensor scheduling, in: Proceedings of SPIE Signal Processing, Sensor Fusion, and Target Recognition V, volume 2755, pp. 304–312.
- [12] K. J. DeMars, I. I. Hussein, C. Fruh, M. K. Jah, R. Scott Erwin, Multiple-object space surveillance tracking using finite-set statistics, *Journal of Guidance, Control, and Dynamics* (2015) 1–16.
- [13] T. A. Hobson, Sensor Management for Enhanced Catalogue Maintenance of Resident Space Objects, Ph.D. thesis, University of Queensland Australia, 2014.
- [14] A. D. Jaunzemis, M. V. Mathew, M. J. Holzinger, Control cost and Mahalanobis distance binary hypothesis testing for spacecraft maneuver detection, *Journal of Guidance, Control, and Dynamics* 39 (2016) 2058–2072.
- [15] R. D. Coder, C. J. Wetterer, K. M. Hamada, M. J. Holzinger, M. K. Jah, Inferring active control mode of the Hubble space telescope using unresolved imagery, *Journal of Guidance, Control, and Dynamics* (accepted) (2017).
- [16] I. I. Hussein, K. J. DeMars, C. Fruh, M. K. Jah, R. S. Erwin, An AEGIS FISST algorithm for multiple object tracking in space situational awareness, in: AIAA/AAS Astrodynamics Specialist Conference, Guidance, Navigation and Control and Co-located Conferences.
- [17] M. J. Holzinger, D. J. Scheeres, K. T. Alfriend, Object correlation, maneuver detection, and characterization using control distance metric, *Journal of Guidance Navigation and Control* 35 (2012) 1312–1325.
- [18] C. Fruh, M. Jah, Initial taxonomy and classification scheme for artificial space objects based on ancestral relation and clustering, in: Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference, Maui, Hawaii.
- [19] G. Shafer, *A Mathematical Theory of Evidence*, Princeton University Press, 1976.
- [20] G. Shafer, R. Logan, Implementing Dempster’s rule for hierarchical evidence, *Artificial Intelligence* 33 (1987) 271–298.
- [21] A. P. Dempster, The Dempster-Shafer calculus for statisticians, *International Journal of Approximate Reasoning* 48 (2008) 365–377.
- [22] T. Denoeux, A k-nearest neighbor classification rule based on Dempster-Shafer theory, *IEEE Transactions on System, Man, and Cybernetics* 25 (1995) 804–813.

- [23] S. L. Hegarat-Masclé, I. Blouch, D. Vidal-Madjar, Application of Dempster-Shafer evidence theory to unsupervised classification in multisource remote sensing, *IEEE Transactions on Geoscience and Remote Sensing* 35 (1997) 1018–1031.
- [24] C. R. Parikh, M. J. Pont, N. B. Jones, Application of Dempster-Shafer theory in condition monitoring applications: A case study, *Pattern Recognition Letters* 22 (2001) 777–785.
- [25] B.-S. Yang, K. J. Kim, Application of Dempster-Shafer theory in fault diagnosis of induction motors using vibration and current signals, *Mechanical Systems and Signal Processing* 20 (2006) 403–420.
- [26] W. F. Caselton, W. Luo, Decision making with imprecise probabilities: Dempster-Shafer theory and application., *Water Resources Research* 28 (1992) 3071–3083.
- [27] D. Aleva, J. McCracken, JSpOC Cognitive Task Analysis, Technical Report, Air Force Research Lab, Wright-Patterson Air Force Base, OH, 2009.
- [28] K. J. Vicente, *Cognitive Work Analysis: Toward Safe, Productive, and Healthy Computer-Based Framework*, CRC Press, 1999.
- [29] C. M. Burns, J. Hajdukiewicz, *Ecological Interface Design*, CRC Press, 2004.
- [30] K. J. Vicente, Ecological interface design: Progress and challenges, *Human Factors* 44 (2002) 62–78.
- [31] K. J. Vicente, J. Rasmussen, Ecological interface design: Theoretical foundations, *IEEE Transactions on Systems, Man, and Cybernetics* 22 (1992) 589–606.
- [32] J. C. Zingarelli, E. Pearce, R. Lambour, T. Blake, C. J. R. Petersen, S. Cain, Improving the space surveillance telescope's performance using multi-hypothesis testing, *The Astronomical Journal* 147 (2014) 111–124.
- [33] T. Hardy, S. Cain, J. Jeon, T. Blake, Improving space domain awareness through unequal-cost multiple hypothesis testing in the space surveillance telescope, *Applied Optics* 54 (2015) 5481–5494.
- [34] T. S. Murphy, M. J. Holzinger, B. Flewelling, Space object detection in images using matched filter bank and Bayesian update, *Journal of Guidance, Control, and Dynamics* 40 (2017) 497–509.
- [35] B. R. Cobb, P. P. Shenoy, On the plausibility transformation method for translating belief function models to probability models, *International Journal of Approximate Reasoning* 41 (2006) 314–330.
- [36] D. Ellsberg, Risk, ambiguity, and the Savage axioms, *The Quarterly Journal of Economics* 75 (1961) 643–669.
- [37] R. Jirousek, P. P. Shenoy, A new definition of entropy of belief functions in the Dempster-Shafer theory, *International Journal of Approximate Reasoning*, in press (2017).
- [38] P. Smets, R. Kennes, The transferable belief model, *Artificial Intelligence* 66 (1994) 191–234.
- [39] G. Shafer, The combination of evidence, *International Journal on Intelligent Systems* 1 (1986) 155–179.
- [40] N. Kami, Integrated formulation of the theory of belief functions from a linear algebraic perspective, in: *IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*, Kongresshaus Baden-Baden, Germany.
- [41] R. R. Yager, Arithmetic and other operations on Dempster-Shafer structures, *International Journal of Man-Machine Studies* 25 (1986) 357–366.
- [42] R. D. Coder, M. J. Holzinger, Multi-objective design of optical systems for space situational awareness, *Acta Astronautica* (2016).
- [43] G. Shafer, Belief functions and parametric models, *Journal of the Royal Statistical Society. Series B* (1982) 322–352.