



Evidence gathering for hypothesis resolution using judicial evidential reasoning

A.D. Jaunzemis^{a,*}, M.J. Holzinger^a, M.W. Chan^b, P.P. Shenoy^c

^a Georgia Institute of Technology, Atlanta, GA 30332, USA

^b Lockheed Martin Space, Sunnyvale, CA 94089, USA

^c University of Kansas, Lawrence, KS 66045, USA

ARTICLE INFO

Keywords:

Dempster–Shafer theory
Hypothesis resolution
Confirmation bias
Ambiguity aversion
Decision support

ABSTRACT

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 rigorous framework for gathering evidence to resolve hypotheses notwithstanding 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 a resolution. The proposed Judicial Evidential Reasoning framework encodes decision-maker questions as rigorously testable hypotheses to be interrogated through evidence-gathering actions. Dempster–Shafer theory is applied to model hypothesis knowledge and quantify ambiguity, and an equal-effort heuristic is proposed to balance time-efficiency and impartiality. Adversarial optimization techniques are used to make many-hypothesis resolution computationally tractable. This work includes derivation of the generalized formulation, computational tractability considerations for improved performance, several illustrative examples, and application to a space situational awareness sensor network tasking scenario. The results show strong hypothesis resolution and robustness to fixation due to poor prior evidence.

1. Introduction

Endsley [1] 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”. Frequently, decision-making occurs in environments where there is insufficient time to gather all available evidence before a decision must be rendered, requiring efficient processes for prioritizing between candidate action sequences. The evidence-gathering or scheduling problem addresses how to obtain, process, and utilize evidence to improve understanding of the state of the environment, as in the use of a sensor network [2]. This is often a high-dimensional, multi-objective, mixed-integer, non-linear optimization problem, so many approaches focus on tractable sub-problems (e.g. single objectives, limited targets, limited actors).

Allocating resources to gather evidence in support of decision-making is an active area of research, and “important and challenging problem in the field of multi-agent systems” [3]. In sensor networks, evidence-gathering often takes the form of target tracking and estimation, such as multi-target tracking (MTT) approaches [4]. Multiple

hypothesis tracking (MHT) is a widely-accepted approach for data association tasks [5], such as maintaining a catalog of space objects [6,7] and associating uncorrelated tracks with a catalog [8]. However, MHT hypotheses are often limited in their form to data association-related questions: e.g. does this new observation belong to an existing tracked object? These types of questions can often be expressed in terms of a state estimate, allowing uncertainty in that state estimate to be reduced through a number of MTT approaches. However, in realistic, complex decision-making scenarios, it is desirable to represent a wider variety of hypotheses (e.g. what is the active mode of this object?) that are not readily related to a state estimate or state uncertainty. This creates a need for the development of generalized hypothesis-resolution approaches that use a rigorous formulation allowing for hypothesis interrogation by gathered evidence [9,10].

Gathered evidence must be fused into a coherent understanding of the environment via association, correlation, and combination [2]. In classical Bayesian approaches, evidence is used to form deterministic probabilities placed on event hypotheses. However, in complex decision-making contexts with uncertainty, evidence also carries ignorance or ambiguity. For this reason, evidential reasoning methods, such as

* Corresponding author.

E-mail addresses: adjaunzemis@gatech.edu (A.D. Jaunzemis), holzinger@ae.gatech.edu (M.J. Holzinger), moses.w.chan@lmco.com (M.W. Chan), pshenoy@ku.edu (P.P. Shenoy).

<https://doi.org/10.1016/j.inffus.2018.09.010>

Received 9 April 2018; Received in revised form 27 July 2018; Accepted 16 September 2018

Available online 17 September 2018

1566-2535/© 2018 Elsevier B.V. All rights reserved.

Dempster–Shafer theory, quantify ambiguity, leading to more realistic modeling of human analyst processes [11–13]. This motivates the use of evidential reasoning approaches, such as Dempster–Shafer theory, to quantify ambiguity and realistically model decision-making processes [11–13]. Dempster–Shafer theory has gained significant traction in various applications, including classification [14,15], monitoring and fault detection [16,17], and decision-making [18].

Another concern in evidence-gathering is confirmation bias, a preferential tendency to gather evidence that confirms prior beliefs [19]. In regimes with uncertainty and ambiguity, this effect also applies by interpreting ambiguous evidence in favor of prior beliefs. Appropriate hypothesis resolution should efficiently and conclusively confirm or refute each proposition while avoiding fixation based on prior information, which may be plagued with uncertainty or ambiguity. Studies have shown that decision-makers demonstrate several biases in decisions involving probability judgment [19–22], so decision-makers must be confident that the evidence-based hypothesis resolution is strong and impartial before declaring a resolution.

The contributions of this work are as follows: (1) a generalized evidence-gathering framework for hypothesis-resolution, (2) the application of evidential reasoning to quantify hypothesis ignorance, (3) a technique for mitigating confirmation bias in action sequence selection, and (4) a computationally tractable approach to the evidence-gathering problem using adversarial optimization techniques.

The remainder of this paper is organized as follows: Section 2 introduces background theory relevant to the theoretical contributions in Section 3, Section 4 presents illustrative examples for application of the framework, Section 5 applies the framework to a simulated space situational awareness problem and Section 6 provides concluding remarks.

2. Background

This section introduces background material relevant to the theoretical developments in the following sections. The focus of this section is introducing the concept of ambiguity aversion, describing its observed impact on decision-making, and introducing the Dempster–Shafer theory of belief functions as a tool to address this phenomenon.

2.1. Ambiguity aversion

Multiple methodologies exist for modeling and reasoning in uncertain domains to provide graphical and numerical representations of uncertainty [23]. 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 [23].

It has been shown that human decision-makers overwhelmingly prefer known risks to unknown risks, making ambiguity a major concern in modeling knowledge states [22]. Ellsberg’s paradox, re-stated here, is a well-known example that violates Savage’s theory of subjective expected utility [22]. 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.

2.2. Dempster–Shafer theory

Dempster–Shafer theory, based on Dempster’s work on probability intervals [24], is considered more expressive than probability theory in representing ambiguity or ignorance [25]. 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 [11], though some alternative formulations such as the Transferable Belief Model allow for relaxation of the latter constraint [26]. Elements of 2^Ω , the set of all subsets of Ω , are referred to as propositions.

2.2.1. Belief functions

A basic probability assignment (bpa) m , as defined in Eq. (1), maps a belief mass to each possible proposition:

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

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

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

Notationally, $\{\theta_1, \theta_2\}$ is equivalent to $\{\theta_1\} \cup \{\theta_2\}$, the disjunctive combination of propositions θ_1 and θ_2 , or “ θ_1 or θ_2 .” The constraint in Eq. (2) enforces the mutually exclusive and collectively exhaustive properties as the belief masses must sum to one, while the constraint in Eq. (3) is similar to Kolmogorov’s axiom of zero probability for the empty set.

Using bpas, 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 [27]. The belief in, or support for, proposition $A \in 2^\Omega$ is defined in Eq. (4) as the sum of belief masses attributed to A and its subsets.

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

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

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

where $\neg A = \Omega \setminus A$ is the negation of A , or “not A .” The representations of belief mass m , belief Bel , and plausibility Pl are all interchangeable via the above linear relationships [28].

2.2.2. Combination

Various bpa combination rules have been developed to fuse evidence from multiple sources into one bpa [29]. The most common combination rule is Dempster's conjunctive rule [27], defined in Eq. (6):

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

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

The normalization constant K , defined in Eq. (7), quantifies the conflict between the two bpas.

2.2.3. Decision-making

While the ability to represent ambiguity in belief functions is useful for accurately representing knowledge states, a key criticism is that the theory of belief functions lacks a coherent decision theory [23]. 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 [26], but it has been argued that the pignistic transformation may not be consistent with Dempster's rule of combination [23]. An alternative method, the plausibility transformation [23], is defined in Eq. (8):

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

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

Note that the normalization constant K in (9) is different from the normalization constant for Dempster's conjunctive rule in Eq. (7). The plausibility transformation is consistent with Dempster's rule, particularly in situations where pignistic probability is inconsistent [23].

Another important concept in both probabilistic and evidential reasoning is entropy as an information content measure. For Dempster–Shafer theory, multiple definitions of entropy have been proposed, many of which are summarized by Jirousek and Shenoy [25]. 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 bpa 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. Numerous other definitions of entropy also exist to address the related concept of hypothesis uncertainty with difference emphases and properties. Jirousek and Shenoy describe a set of desirable traits for any bpa entropy definition and compare a number of existing definitions against this list [25].

One useful property of J–S entropy, in contrast to similar entropy definitions such as Deng entropy [30], is that maximum entropy is only attained by a vacuous bpa, which is the bpa 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 [25]. Recalling Ellsberg's paradox, the first urn is an equally-likely belief structure and the second urn is a vacuous belief structure:

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

The Shannon entropy, Dubois–Prade entropy, and J–S entropy for these belief structures are shown in Table 1. 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




Fig. 1. Predictive and reactive evidence-gathering.

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.

Importantly, J–S entropy scales both the conflict and non-specificity portions equally, allowing the use of this single objective without the need for scaling parameters. Therefore, minimizing J–S entropy can be used as a reliable and consistent metric for a strong hypothesis resolution. For further discussion and comparison of entropy definitions, refer to Deng [30] or Jirousek and Shenoy [25].

3. Theory

The following section builds upon the background material from the previous section to motivate the development of Judicial Evidential Reasoning (JER). The JER approach hinges upon three primary considerations: hypothesis abstraction, ambiguity aversion, and unbiased hypothesis resolution. The theory is developed as a general evidence-gathering approach applicable to a wide variety of hypothesis-resolution tasks, as shown by the medical diagnosis examples in Section 4 and the sensor network tasking application in Section 5.

3.1. Hypothesis abstraction

Many evidence-gathering approaches (e.g. sensor network tasking) operate on maintaining a low overall uncertainty (e.g. information-maximum); however, it may not be readily apparent to a decision-maker how reducing state uncertainty by a certain amount affects situation awareness or answers decision-making questions. This motivates an approach that encodes decision-making priorities as hypotheses that can be interrogated by evidence-gathering actions.

Hypothesis-driven approaches enable a predictive mode of evidence-gathering designed to answer specific questions, using prior knowledge of relevant hypotheses to estimate information-gain from potential courses of action and propose actions that are predicted to resolve the hypotheses. This is fundamentally different from reactive approaches, where the gathered information is used to form hypotheses *a posteriori* about what caused the observed behavior. This relationship between hypotheses and information in predictive and reactive evidence-gathering is illustrated in Fig. 1.

Re-framing evidence-gathering in terms of hypotheses supports human decision-making strengths in abstract-level cognitive tasks required for objective prioritization and goal-adjustment [31]. Forcing an operator to switch between different levels of the abstraction, effectively approaching the problem at multiple different levels of detail, leads to increased frustration and workload and decreased situation awareness [10]. Designing a decision-support system that directly conveys hypothesis resolution information ensures that the human decision-maker spends more time on strategic cognitive tasks.

3.2. General evidence-gathering problem definition

Consider a set of hypotheses and a set of actors tasked with gathering evidence to resolve these hypotheses over a given T -step time horizon,

Table 1
Ellsberg's paradox belief structures and entropy.

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

from t_k to t_{k+T} . The finite set of hypotheses under consideration can be represented as $\Omega = \{\Omega_1, \dots, \Omega_n\}$, where Ω_i is the frame of discernment for the i th hypothesis and $|\Omega| = n \in \mathbb{Z}^+$ is the number of hypotheses. Recall that each hypothesis frame of discernment $\Omega_i \in \Omega$ contains a set of mutually exclusive and collectively exhaustive propositions for resolution of that hypothesis.

At time t_k , define the actions available to the s th actor as the finite set $\mathbb{A}_{s,k}$. Under the assumption that each actor can only perform one action at a given time t_k , the available action sets for all m actors in the network at time t_k are described through the Cartesian product:

$$\mathbb{A}_k = \mathbb{A}_{1,k} \times \dots \times \mathbb{A}_{m,k} \quad (10)$$

where an action set $\mathcal{A}_k \in \mathbb{A}_k$ denotes a valid set of m actions at time t_k and $A_{s,k} \in \mathcal{A}_k$ is a valid action for actor s from that action set.

Define an actor's sequence of actions over the time horizon t_{k+1} to t_{k+T} as the following ordered list (or T -tuple):

$$\mathcal{A}_{s,1:T} = (A_{s,1}, \dots, A_{s,T}), \quad s = 1, \dots, m \quad (11)$$

Similarly, define a set of action sequences for all actors as the finite set:

$$\mathcal{A}_{1:T} = \{\mathcal{A}_{1,1:T}, \dots, \mathcal{A}_{m,1:T}\} \quad (12)$$

This set contains an action sequence for each of the m actors in the network and thus fully defines the actions taken by the network over the time horizon. Furthermore, the set of all valid sets of action sequences (all valid combinations of action sequences) is also represented by a Cartesian product:

$$\mathbb{A}_{1:T} = \mathcal{A}_{1:T} \times \dots \times \mathcal{A}_{1:T} \quad (13)$$

The goal is to select the set of action sequences that minimizes a to-be-defined cost function at the end of the T -step receding time horizon. Generically, this cost function may be represented as follows:

$$J_T : (\Omega, \mathcal{W}; \mathcal{A}_{1:T}) \mapsto \mathbb{R} \quad (14)$$

where $\mathcal{W}$ is a user-defined set of weights such that $w_i \in \mathcal{W}$ quantifies the priority of hypothesis Ω_i relative to the other hypotheses in Ω , and T indicates that the cost function is evaluated at the end of the time horizon, time t_{k+T} . It stands to reason that some hypotheses will be more important to decision-makers than others, so this weighting is considered a user-defined (potentially time-varying) parameter. It is not subject to optimization in this study but is instead treated as a tunable parameter.

Therefore, the generic hypothesis-based evidence-gathering optimization problem is:

$$\mathcal{A}_{1:T}^* = \arg \min_{\mathcal{A}_{1:T} \in \mathbb{A}_{1:T}} J_T(\Omega, \mathcal{W}; \mathcal{A}_{1:T}) \quad (15)$$

In other words, the optimal set of action sequences minimizes the cost function J_T , evaluated at time t_{k+T} subject to the evidence from each action $A_{s,t} \in \mathcal{A}_{s,1:T}$ in each action sequence $\mathcal{A}_{s,1:T} \in \mathcal{A}_{1:T}$ for each actor $s = 1, \dots, m$. In the following sections, a specific cost function is developed based on reaching strong (unambiguous and unbiased) hypothesis resolutions.

3.3. Evidence-gathering for hypothesis entropy reduction

Hypothesis resolution refers to the goal of determining which proposition is true from the set of propositions in the frame of discernment. Recall that Jirousek–Shenoy (J–S) entropy [25] quantifies both conflict and non-specificity in hypothesis knowledge, providing an apt minimization

objective for strong hypothesis resolution. In particular, the conflict and non-specificity portions of J–S entropy have equal scaling, allowing for the use of this single objective in the measurement of both phenomena without need for scaling parameters. Many different measures of bpa entropy have been presented in recent years, including some which also incorporate both conflict and non-specificity, such as Deng entropy [30]. J–S entropy has the additional advantage of reaching maximum entropy only with a vacuous bpa, providing a clear indication of when conflict and non-specificity are both at a maximum. As this paper focuses on the development of a hypothesis-resolution approach as a whole, comparisons between different implementations of hypothesis entropy will be left for future work. J–S entropy quantifies (in equal scaling) both conflict and non-specificity and ensures that the vacuous bpa receives maximum entropy, which satisfies the desired entropy properties for this hypothesis-resolution task.

At a given time t_k , each candidate action $A \in \mathcal{A}_k$ gathers evidence that may be used to resolve hypotheses. Denote the total amount of evidence gathered through action set $\mathcal{A}_k$ as p , noting that a single action may gather multiple distinct pieces of evidence or may gather no evidence, restricting p to the non-negative integers. The hypothesis-resolution contribution of a given piece of evidence is represented by the bpa:

$$m_{i,j,k} : 2^{\Omega_i} \mapsto [0, 1] \quad (16)$$

where the subscript i indicates that this bpa is related to hypothesis Ω_i , the subscript $j = 1, \dots, p$ refers to the piece of evidence relevant to this bpa, and the subscript k indicates the evidence is gathered at time t_k . The bpas for all p pieces of evidence can be fused using Dempster's rule to arrive at a hypothesis update bpa:

$$\tilde{m}_{i,k} = \bigoplus_{j=1}^p m_{i,j,k} \quad (17)$$

Recall that Dempster's rule is associative and commutative, meaning the combination can be done sequentially and order doesn't matter [32]. However, Dempster's rule is not idempotent so the pieces of evidence being combined must be independent to avoid artificially inflating the effect of a particular piece of evidence. If a particular piece of evidence j gathered at time t_k does not contribute to hypothesis Ω_i , then $m_{i,j,k}$ is simply the vacuous bpa, ensuring that each term in the summation is defined.

Therefore, the resulting knowledge state for hypothesis Ω_i , incorporating all evidence from time from t_0 to t_k is denoted as:

$$m_{i,k}^+ = m_{i,k}^- \oplus \tilde{m}_{i,k} \quad (18)$$

where $m_{i,k}^+$ is the a posteriori knowledge state and $m_{i,k}^-$ is the a priori knowledge state for hypothesis Ω_i at time t_k based on all evidence gathered prior to t_k .

Note that, while this approach implements evidence fusion, combination rules (i.e. Dempster's rule), and hypothesis entropy definitions, improving upon evidence combination rules or entropy definitions is not a focus of this work. Instead, the goal of this work is to develop an improved method for evidence search and resource allocation.

3.3.1. Normalized Jirousek–Shenoy entropy

The resolution of hypothesis Ω_i based on bpa m_i , as measured through J–S entropy [25], is defined as:

$$H_{JS}(m_i) = \left(\sum_{x \in \Omega_i} \text{Pl}_{p_{m_i}}(x) \log_2 \left(\frac{1}{\text{Pl}_{p_{m_i}}(x)} \right) \right) + \left(\sum_{A \in 2^{\Omega_i}} m_i(A) \log_2(|A|) \right) \quad (19)$$

where the first summation term, related to Shannon entropy, quantifies conflict and the second summation term, called Dubois–Prade entropy, quantifies non-specificity. [Theorem 3.1](#) shows that J–S entropy is on the scale $[0, 2\log_2(|\Omega_i|)]$ [\[25\]](#).

Theorem 3.1 (Maximum Entropy). *Consider a bpa m for discrete random variable X with frame of discernment $\Omega_X = \{x_1, \dots, x_n\}$. The maximum value of J–S entropy for m is*

$$H_{JS}(m) = 2\log_2(n) \quad (20)$$

Proof. The maximum entropy principle [\[25\]](#) states that the maximum value of entropy is attained by the vacuous bpa. Therefore, assume m is vacuous:

$$m(\Omega_X) = 1, m(x) = 0 \forall x \in 2^{\Omega_X} \setminus \Omega_X \quad (21)$$

The plausibility transformation yields the following plausibility probability function:

$$Pl_{P_m}(x) = \frac{1}{n} \forall x \in \Omega_X \quad (22)$$

Computing the Shannon entropy:

$$\sum_{x \in \Omega_i} Pl_{P_{m_i}}(x) \log_2 \left(\frac{1}{Pl_{P_{m_i}}(x)} \right) = n \left(\frac{1}{n} \log_2(n) \right) = \log_2(n) \quad (23)$$

Similarly, computing the Dubois–Prade entropy for the vacuous bpa m :

$$\sum_{A \in 2^{\Omega_i}} m_i(A) \log_2(|A|) = \log_2(n) \quad (24)$$

because $m(A)$ is non-zero only for $A = \Omega_X$, in which case the cardinality $|A| = |\Omega_X| = n$. Maximum J–S entropy is therefore the sum: $H_{JS}(m) = 2\log_2(n)$. $\square$

To use entropy as a cost function while accounting for hypotheses with different numbers of propositions, the normalized J–S entropy is defined as:

$$\tilde{H}_{JS}(m_i) = \frac{H_{JS}(m_i)}{2\log_2(|\Omega_i|)} \quad (25)$$

where m_i is the bpa representing knowledge of hypothesis Ω_i .

3.3.2. Optimization formulation

To accomplish the goal of minimizing hypothesis conflict and non-specificity, the normalized entropy defined in [Eq. \(25\)](#) is employed as the cost function to further specify the optimization problem in [Eq. \(15\)](#).

$$\mathcal{A}_{1:T}^* = \arg \min_{\mathcal{A}_{1:T} \in \mathbb{A}_{1:T}} \sum_{i=1}^{|\Omega|} w_i \tilde{H}_{JS}(\hat{m}_{i,T}) \quad (26)$$

where $w_i \in \mathcal{W}$ are the hypothesis weights used to denote relative priorities such that $\sum_i w_i = 1$, and $\hat{m}_{i,T}$ is the estimated bpa for hypothesis Ω_i at the end of the time horizon t_{k+T} . In this work, the weights w_i are considered tunable parameters that the operator may modify to reflect relative priorities between the hypotheses (e.g. some hypotheses may be more time-critical than others). Selecting an optimal set of weights is a challenging multi-objective optimization problem of its own that is not within the scope of this development; instead, this work shows that, given a set of weights determined by the operator, the JER algorithm is able to select an optimal set of action sequences. This optimal set of action sequences, $\mathcal{A}_{1:T}^*$, are actions estimated to gather evidence that minimizes conflict and non-specificity in user-prioritized hypotheses.

3.3.3. Computational complexity

The general formulation in [Eq. \(26\)](#) suffers from a number of practical issues in implementation. Most notably, the number of action sequences to evaluate can quickly preclude brute-force evaluation of all possible action sequences over the time horizon. Computational complexity of a brute-force approach to this optimization problem scales

with the number of hypotheses $|\Omega|$, the number of sensors m , the number of valid actions for each sensor n_m , and the time horizon T :

$$\mathcal{O} \left(\prod_{t=1}^T \left(\prod_{j=1}^m n_{s,t} \right) \right) \quad (27)$$

where m is the number of actors and $n_{s,t}$ is the number of valid actions for the s th actor at time t_{k+t} . The upper bound on this complexity is found by defining the worst-case number of valid actions as $n = \max(n_{s,t} \mid s = 1, \dots, m; t = 1, \dots, T)$, yielding a worst-case computational complexity:

$$\mathcal{O}((n^m)^T) \quad (28)$$

As expected, computational complexity for a brute-force approach scales exponentially with the number of valid actions, the number of actors, and the length of the time horizon. Depending on the resources required to estimate the hypothesis resolution after a set of action sequences, this algorithm can become computationally restrictive, motivating several complexity mitigations.

3.4. Implementation considerations

This section modifies the general optimization approach in [Eq. \(26\)](#) to arrive at a computationally tractable solution by decomposing the problem into individual sub-problems and applying adversarial optimization techniques to reduce the number of action sequence evaluations. An additional concern with the entropy-reduction algorithm is the effect of evidence ambiguity and evidence-gathering bias induced by prior information. Adversarial optimization is applied to reduce the number of action sequence evaluations and combat confirmation bias.

3.4.1. Unbiased hypothesis resolution

Confirmation bias is a cognitive phenomenon where prior belief causes fixation on a particular proposition, causing the human to favor evidence that confirms prior beliefs and overlook conflicting evidence [\[19\]](#). In regimes with uncertainty and ambiguity, this effect also applies by interpreting ambiguous evidence in favor of prior beliefs. Similar to human cognitive fixation, socio-technical systems might also exhibit confirmation bias. For instance, a most-probable-first evidence-gathering approach would prioritize actions estimated to gather further evidence to confirm prior knowledge. However, spurious detections or false alarms may lead to increased belief in the incorrect proposition. In this way, prior information has the potential to skew future evidence-gathering actions, so technological fixation may be induced by measurement noise, sensor bias, or other sources of uncertainty.

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 approach would focus future actions on confirming $\{x_1\}$, while ignoring the (much larger) ignorance in the estimated proposition. If the true resolution of this hypothesis is actually $\{x_2\}$, evidence gathered from tasking on the incorrect proposition ($\{x_1\}$) may be vacuous, causing the knowledge state stagnate. In this case, the most-probable-first approach stalls as no further evidence is admitted to increase belief in $\{x_2\}$ and change the proposed tasking.

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. 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. Quantifying confirmation bias is an active area of research, with cognitive sciences researchers using various measures comparing selection of supporting versus refuting evidence [\[19,21\]](#). One such measurement is the difference between numbers of selected supporting and refuting evidence elements [\[20\]](#), meaning an unbiased sequence of actions selects equal numbers of supporting and refuting elements.

The proposed approach employs a related heuristic, a principle of equal effort that distributes resources (e.g. actions, time, money) evenly amongst propositions. An apt analogy for this heuristic is the fair trial system, wherein the defense and prosecution are given equal opportunity to present the strongest evidence to confirm or refute a hypothesis. Similarly, the proposed framework employs a pair of agents for each proposition, advocate and critic, which alternate action turns to allow equal opportunity for gathering supporting or refuting evidence, respectively. Due to strong parallels to the fair trial system, the proposed framework is called Judicial Evidential Reasoning (JER).

Application of this alternating-turns heuristic encourages resolution guided by evidence, not prior beliefs, biases, or ambiguity. In the event of multiple competing resources, the principle of equal effort creates an additional multi-objective optimization and uniqueness of the solution using this heuristic is not guaranteed. However, improved measures for confirmation bias are an area for future research and could extend the JER approach by altering the agent-pair action ordering.

3.4.2. Sub-problem definition

The primary intuition that allows decomposition of the entropy-reduction approach in Eq. (26) is that not all sensor actions contribute evidence related to all hypotheses. The sub-problems can be solved independently (and in parallel), resulting in $|\Omega|$ sub-problem action sequence sets that must be combined into a single optimal set of action sequences.

Consider one of the hypotheses $\Omega_i \in \Omega$ and the subset of valid actions relevant to that hypothesis as $\mathbb{A}_{s,k,i}$:

$$\mathbb{A}_{s,k,i} \subseteq \mathbb{A}_{s,k} \quad (29)$$

where $\mathbb{A}_{s,k}$ are all the valid actions for sensor $s = 1, \dots, m$. Similarly, the action sequences relevant to hypothesis Ω_i over the time horizon t_k to t_{k+T_i} are denoted

$$\mathbb{A}_{1:T_i,i} \subseteq \mathbb{A}_{1:T_i} \quad (30)$$

By definition, $|\mathbb{A}_{s,k,i}| \leq |\mathbb{A}_{s,k}|$ and $|\mathbb{A}_{1:T_i,i}| \leq |\mathbb{A}_{1:T_i}|$. Note that the time horizon T_i is allowed to be different for each hypothesis since, in operation, not all hypotheses need to have the same optimization horizon.

The sub-problem optimization objective is first represented using a generic cost function specific to each hypothesis:

$$J_{T_i,i} : (\Omega_i; \mathbb{A}_{1:T_i,i}) \mapsto \mathbb{R} \quad (31)$$

Note that w_i is not relevant to this portion of the optimization as the sub-problems are being solved independently, but it will play a role in the combination of the sub-problem sequences. The sub-problem optimization problem is defined as:

$$\mathcal{A}_{1:T_i,i}^* = \arg \min_{\mathcal{A}_{1:T_i,i} \in \mathbb{A}_{1:T_i,i}} J_{T_i,i}(\Omega_i; \mathcal{A}_{1:T_i,i}) \quad (32)$$

This sub-problem decomposition approach allows for parallel computation of action sequence sets for each agent-pair. However, if the entropy-reduction cost function is employed as in Eq. (26), the same concerns related to confirmation bias will arise: an incorrect prior induces actions against the incorrect proposition, leading to weak or vacuous evidence and weak hypothesis resolution. Therefore, a different optimization approach is employed for the sub-problems while entropy-minimization is reserved for the combination of the sub-problem solutions.

3.4.3. Combating confirmation bias

Adversarial optimization techniques are employed to reduce confirmation bias, similar to the opposing counsel in the judicial system. Approaches such as minimax optimization have been heavily applied in game theory for turn-based, zero-sum games such as Chess and GO. In minimax optimization, an agent plans its actions with the knowledge that the opposing agent will select actions toward the opposite goal. In light of this conflict, both agents attempt to minimize potential loss in

a worst-case scenario. Conversely, for a maximizing objective, maximin optimization represents agents maximizing the minimum gain from a sequence of actions. The adversarial minimax approach ensures that the evidence gathered equitably presents available hypothesis resolution information while remaining impartial to the hypothesis resolution result.

Consider a single hypothesis from the set of considered hypotheses at time t_k : $\Omega_i \in \Omega$. Each proposition must be either conclusively confirmed or refuted with evidence, so each proposition is assigned a pair of JER agents. Therefore, for hypothesis Ω_i there are $|\Omega_i|$ alternating JER agent-pairs. When the advocate agent is active, its goal is to maximize belief in the proposition $\{\theta\}$, accomplished using maximin optimization with the plausibility probability transformation:

$$\mathcal{A}_{1:T_i,i}^* = \arg \max_{\mathcal{A}_{1:T_i,i} \in \mathbb{A}_{1:T_i,i}} \min_{\theta} \Pr_{\text{pl}}(\theta; m_{i|\mathcal{A}_{1:T_i,i}}) \quad (33)$$

where $m_{i|\mathcal{A}_{1:T_i,i}}$ is the estimated bpa 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. The maximum attainable value for this objective is 1 when proposition $\{\theta\}$ has full belief, and the minimum attainable value for this objective is 0 when proposition $\{\neg\theta\}$ has full belief. When the critic agent is active, its goal is to maximize belief in the alternative proposition $\{\neg\theta\}$ or equivalently minimize belief in the null proposition $\{\theta\}$. Therefore, the formulation simply flips to a minimax optimization:

$$\mathcal{A}_{1:T_i,i}^* = \arg \min_{\mathcal{A}_{1:T_i,i} \in \mathbb{A}_{1:T_i,i}} \max_{\theta} \Pr_{\text{pl}}(\theta; m_{i|\mathcal{A}_{1:T_i,i}}) \quad (34)$$

The result of the JER agent-pair schedule optimization is a minimax-optimal action sequence for each agent-pair. In the next section, these sub-problem action sequences are combined to arrive at a single optimal schedule. If an agent-pair's action is selected in the final schedule for this iteration, that agent-pair flips its active agent for the next time step.

3.4.4. Resolving combined schedule incongruity

After determining optimal schedules for each agent-pair, the schedules must be combined into a single schedule. Depending on the hypotheses, it is possible or even likely that two or more agent-pairs will require the same actor for different actions. These incongruities are resolved by choosing the actions that lead to the strongest hypothesis resolution as measured by entropy.

Using the set of actions from all sub-problem optimal sequences $\mathcal{A}_{1:T,i}^*$, all possible combinations of these actions are used to form candidate congruous action sequences. The combination schedules are evaluated up to the longest time horizon. The evaluation criterion for selecting the optimal combined schedule is the weighted-sum of entropy:

$$\mathcal{A}_{1:T}^* = \arg \min_{\mathcal{A}_{1:T} \in \mathbb{A}_{1:T}^*} \sum_{i=1}^{|\Omega|} w_i \tilde{H}_{JS}(m_{i|\mathcal{A}_{1:T}}) \quad (35)$$

where w_i is the weighting for the i th hypothesis, and $\tilde{H}_{JS}$ is the normalized J-S entropy as defined in Eq. (25). Since J-S entropy quantifies both conflict and non-specificity, and the weighting parameters encode decision-maker priorities, the resulting action sequence $\mathcal{A}_{1:T}^*$ is the action sequence with the strongest priority-weighted resolution.

At worst case, this is the same as a brute-force re-evaluation, but this would require all hypotheses to 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 at least some hypotheses are distinct enough to have different applicable actions, this re-evaluation is much less computationally complex than brute-force.

3.4.5. Efficient minimax optimization

To further reduce the number of action sequences evaluated, the alternating-agent formulation of the sub-problems can be further exploited using adversarial optimization techniques. Combinatorial optimization techniques often employ methods for intelligently exploring




Fig. 3. Judicial evidential reasoning algorithm.

solves its sub-problem using the alternating-agent minimax optimization on the plausibility probability transformation, outlined in Algorithm 2 using a naive minimax algorithm for ease of description. Recall that the alpha-beta pruning enhancement is simply an efficient minimax that reduces the search space, so both the naive and alpha-beta implementations reach the same result. The JER manager combines the sub-problem schedules and determines the optimal combined action sequence $\mathcal{A}^{**}$ using the total weighted entropy objective. Once the optimal action sequence is determined, the active agent-pairs (those whose actions are chosen in the optimal sequence) are flipped for the next iteration.

The following sections apply the JER algorithm described above to illustrative medical diagnosis examples as well as a sensor network tasking scenario.

4. Examples

This section contains illustrative examples of the JER approach using simplified medical diagnosis situations. The first example is intended to illustrate the JER agent-pair sub-problem optimization in detail, and the second example illustrates the combination of multiple JER agent-pairs to form a unified schedule. In each case, the relevant hypotheses and propositions are outlined, available tests are outlined as potential

actions, and the JER approach is applied to determine the sequence of tests ordered. Since this is intended to be an illustrative example, the diagnosis details have been simplified and constraints have been enforced such that not all actions may be taken within the diagnosis window (i.e. time-critical decision-making). Section 5 contains a more detailed and nuanced application of JER to a real-world sensor network tasking problem.

4.1. Case 1: single JER agent-pair

The first example involves a single hypothesis with two competing (mutually exclusive and collectively exhaustive) propositions, yielding the following frame of discernment: $\{\theta, \neg\theta\}$. To simplify notation, define the corresponding propositions as: $A = \{\theta\}$ and $\neg A = \{\neg\theta\}$. Since this is a single hypothesis problem with a binary frame of discernment, only one JER agent-pair is needed, and this example serves to illustrate the JER inner-loop: minimax optimization using plausibility probability.

Assuming no prior information on the correct resolution (i.e. full ignorance), the prior belief assignment is vacuous; this is to say that each proposition has the same belief and plausibility, zero and one respectively, prior to any evidence gathering. Three tests are available to inform this diagnosis, but time and cost constraints limit the number of

Algorithm 2 JER agent-pair evaluation, recursive naive minimax implementation.

```

1:  $t_k$ : current time
2:  $t_T$ : horizon time
3:  $\hat{m}_k^-$ : estimated a priori bpa at time  $t_k$ 
4:  $isMax$ : flag, true if advocate is active at  $t_k$ 
5:  $\mathcal{A}$ : action sequence
6: procedure EVALUATEAGENTPAIR( $t_k, t_T, \hat{m}_k^-, isMax, \mathcal{A}$ )
7:   if  $t_k \geq t_T$  then
8:     return ( $\Pr_{pl}(\{\emptyset\}), \mathcal{A}$ )
9:   else
10:    if  $isMax$  then
11:       $s^* \leftarrow -\inf$ 
12:    else
13:       $s^* \leftarrow +\inf$ 
14:    end if
15:     $\mathcal{A}^* \leftarrow \mathcal{A}$ 
16:     $\mathbb{A}_k \leftarrow$  candidate actions relevant to this agent-pair at  $t_k$ 
17:     $t_{kp} \leftarrow$  next time step
18:    for each action set  $A$  in  $\mathbb{A}_k$  do
19:       $\mathcal{A}[t_k] \leftarrow A$ 
20:       $\hat{m}_A \leftarrow$  estimated bpa from action set  $A$ 
21:       $\hat{m}^+ \leftarrow \hat{m}^- \oplus \hat{m}_A$ 
22:       $(s, \mathcal{A}^+) \leftarrow \text{EvaluateAgentPair}(t_{kp}, t_T, \hat{m}_k^+, isMax, \mathcal{A})$ 
23:      if  $isMax$  then
24:        if  $s > s^*$  then
25:           $s^* \leftarrow s$ 
26:           $\mathcal{A}^*[t_k : t_T] \leftarrow \mathcal{A}^+[t_k : t_T]$ 
27:        end if
28:      else
29:        if  $s < s^*$  then
30:           $s^* \leftarrow s$ 
31:           $\mathcal{A}^*[t_k : t_T] \leftarrow \mathcal{A}^+[t_k : t_T]$ 
32:        end if
33:      end if
34:    end for
35:    return ( $s^*, \mathcal{A}^*$ )
36:  end if
37: end procedure

```



Table 2

Example Case 1: Basic probability assignments for diagnostic tests.

Test #	A	$\neg A$	$A \cup \neg A$
1	0.7	0.0	0.3
2	0.0	0.7	0.3
3	0.3	0.3	0.4

tests to two. Therefore, the goal is to determine which two tests result in a strong-but-unbiased resolution. Table 2 lists basic probability assignments (bpas) for each available test, functions of known statistics on the test such as false alarm rate. Test 1 is a strong indicator of the confirmation of proposition A but does not carry information to negate A (or equivalently, confirm $\neg A$). Similarly, Test 2 is a strong indicator for confirming proposition $\neg A$. Both tests have non-zero probabilities of false alarm, meaning that neither can contribute complete belief to either proposition, resulting in non-zero belief mass attributed to the frame. Test 3 does operate as an indicator of both confirmation and negation of A, but provides weaker evidence toward both.

Fig. 4 shows the tree of all possible evaluations, easily visualized due to the low dimensionality of this example. Each edge of this tree is a potential test (action), and each terminating node denotes the plausibility probability of A as a result of the two actions leading to it. For instance,

**Fig. 4.** Example Case 1: Brute force evaluation tree.**Fig. 5.** Example Case 1: Alpha-beta evaluation tree.

the left cluster of three terminating nodes represents all possible action sequences beginning with Test 1. Traversing down the tree, each successive level alternates the active agent: supporting A or $\neg A$. This results in a two-step minimax optimization on $\Pr_{pl}(A)$. The non-terminating nodes display the chosen node from below based on the active minimax mode at that step (max or min). Therefore, the minimizing agent (supporting $\neg A$) will select the lowest plausibility probability from each cluster of three terminating actions to populate the middle nodes, and the maximizing agent (supporting A) will select the highest plausibility probability from those three middle nodes to determine the selected action sequence (highlighted in blue in the figure).

Following test 1 with another test 1 yields an estimated bpa that may strongly indicate the diagnosis A but does not carry any unique evidence to confirm $\neg A$. Therefore, if $\neg A$ is the correct result, performing test 1 twice would not provide a strong result. Following test 1 with test 2 or test 3 results in estimated belief mass attributed to both A and $\neg A$, resulting in (at least partial) proposition confirmation regardless of the correct (true) result.

Fig. 4 shows that the unbiased solution in this minimax optimization scheme is Test 1 followed by Test 2, resulting in $\Pr_{pl}(A) = \Pr_{pl}(\neg A) = 0.5$. This indicates that both propositions are given equal opportunity since the prior information (vacuous) did not indicate a preference toward either proposition. This result matches intuition that, in the case of vacuous prior information, both strong indicator tests should be run to ensure the true diagnosis is confirmed.

In Fig. 4, all nine possible test sequences are computed in a brute-force manner. However, this is not required as alpha-beta pruning provides a more efficient approach to minimax optimization by eliminating branches of the evaluation tree that need not be searched based on the previously-searched nodes. Fig. 5 demonstrates this approach, reducing the number of sequence evaluations from nine to seven. As expected, alpha-beta pruning efficiently finds the same end-result as naive brute-force minimax using less evaluations.

4.2. Case 2: multiple JER agent-pairs

The second example involves a single hypothesis with three competing (mutually exclusive and collectively exhaustive) propositions, yield-

Table 3
Example Case 2: Basic probability assignments for diagnostic tests.

Test #	A	B	C	$A \cup B$	$A \cup C$	$B \cup C$	$A \cup B \cup C$
1 (Pass)	0.9	0.0	0.0	0.0	0.0	0.0	0.1
1 (Fail)	0.0	0.0	0.0	0.0	0.0	0.9	0.1
2 (Pass)	0.0	0.7	0.0	0.0	0.0	0.0	0.3
2 (Fail)	0.0	0.0	0.0	0.0	0.7	0.0	0.3
3 (Pass)	0.0	0.0	0.5	0.0	0.0	0.0	0.5
3 (Fail)	0.0	0.0	0.0	0.5	0.0	0.0	0.5

ing the following frame of discernment: $\{\theta_1, \theta_2, \theta_3\}$. Once again, for ease of notation, define the corresponding propositions as: $A = \{\theta_1\}$, $B = \{\theta_2\}$, $C = \{\theta_3\}$. Since this case contains a non-binary frame of discernment, three JER agent-pairs are needed (to ensure one advocate for each proposition). Each agent-pair solves its own minimax sub-problem (as in Case 1), and the sub-problem solutions are combined to resolve incongruity between the sub-problem schedules. Therefore, this example serves to illustrate the application of the entropy-minimization objective to resolve schedule incongruity, as well as updating the prior with test results iteratively in receding-horizon optimization.

As before, the prior belief assignment is assumed vacuous: each proposition starts with a belief of zero and plausibility of one, representing full ignorance or ambiguity. Three tests are available to inform this diagnosis, but time and cost constraints limit the number of tests to two. Table 3 lists basic probability assignments (bpas) for each available test, functions of known statistics on the test such as false alarm rate. Each test has two possible outcomes, pass and fail, which affect the knowledge state based on the test's evidence and therefore affect the test bpas. For instance, test 1 is a strong indicator for proposition A, so a pass outcome gives strong belief for A whereas a fail outcome gives strong belief for $B \cup C = \neg A$. In both cases, though, there is still a non-zero chance of false alarm, so some belief mass is still given to the frame as ignorance. Tests 2 and 3 are similar for propositions B and C, respectively, though test 2 is only moderately strong and test 3 is even weaker.

For each iteration, each JER agent-pair must solve its own sub-problem, selecting the two minimax-optimal tests in a similar manner to the previous example case, using the alpha-beta pruning improvement. Each agent operates under the supposition that its desired proposition is correct. In other words, the advocate agent for proposition A estimates that test 1 will return successful while tests 2 and 3 will return failed, attempting to contribute belief to proposition A. On the other hand, the critic agent for proposition A estimates the opposite test results, since it is attempting to reduce plausibility of proposition A.

After the sub-problem test schedules are optimized, all unique test sequences are found from the Cartesian product of these schedules. The estimated bpa result is computed for each under the assumption of a successful test return, and the sequences are ranked according to normalized J-S entropy. The test sequence resulting in the lowest entropy is selected, and its first test action implemented for that iteration. Finally, the agent-pairs that were active during this iteration (i.e. the agent-pairs that requested the selected action) are flipped so that, in the next iteration, the critic agent is active first.

In order to simulate this algorithm in action, one of the propositions must be selected as the “truth” hypothesis resolution, which affects the result of each test. For instance, if proposition A is assumed to be true, then test 1 would pass while tests 2 and 3 would fail. The appropriate evidence bpa is selected from Table 3 based on the test result, and the result after one iteration is computed by combining the prior bpa with this evidence bpa. The procedure above is repeated for the second test iteration in a receding-horizon approach, using the updated bpa as the new prior.

Table 4 shows three different simulated example cases: the first where A is the true resolution, another where B is true, and the third where C is true. In each case, a first test is selected and the vacuous

prior is updated with the appropriate evidence bpa from Table 3 based on the case's assumed true resolution. A second test is then selected based on the updated bpa, and the appropriate evidence bpa is used to further updating the hypothesis knowledge. Table 4 details the hypothesis resolution after each test to illustrate the evolution of probability and entropy.

Since the initial prior is vacuous, the chosen first test in each realization is test 1, matching intuition because test 1 yields the strongest result and therefore minimizes entropy. If proposition A is true, as in the top row of Table 4, test 1 passes and significant belief is already attributed to A. In the second iteration, test 1 is repeated because, in this case, the entropy will not be significantly decreased by the contribution of other (weaker) test results. Even though the critic agent is active for A, test 1 is still the strongest potential belief contribution for $\neg A$, providing the strongest entropy reduction. Instead, test 1 passes again, confirming proposition A and further reducing entropy since A is the probable correct resolution. If either proposition B or C are true, test 1 fails, resulting in significant ambiguity after the first test (as seen in the higher normalized J-S entropy after test 1). In both cases, the next test selected is test 2 because it is the strongest remaining test and proposition A has (nearly) been eliminated. If proposition B is true, test 2 passes and now a significant belief is attributed to B. If proposition C is true, test 2 fails as well, resulting in a less-significant but still distinct belief attributed to C: the fail result in test one indicates B or C and the fail result in test 2 indicates A or C, leaving C as the only logically consistent option. In this case, the belief attributed to A slightly increases after the second test result, but evidence still overwhelmingly indicates proposition C.

Once again, this result matches intuition. Since test 3 has the weakest belief contribution and only two tests may be executed, it is never selected. The correct resolution is determined with high probability through the use of only two tests in each case.

Note that, while this example case did require more test sequence evaluations than a brute-force implementation, this is only because each available test is relevant to each agent-pair (for simplicity). Computational complexity of this evaluation scales exponentially with the number of available actions, so in a low-dimensional scenario such as this example, more available actions may not be an issue. In real-life complex decision-making scenarios with multiple propositions or hypotheses, it is likely that this computational burden can be significantly reduced by only considering relevant actions for each agent-pair.

4.3. Example summary

The two example cases presented illustrate both key components of the JER approach. Case one demonstrates the inner-loop sub-problem resolution using agent-pairs and efficient minimax with the plausibility probability transformation. Case two demonstrates the outer-loop combination of sub-problem schedules and resolution of incongruities using entropy.

While both cases illustrated use a single hypothesis, multiple hypotheses do not significantly change the implementation. New agent-pairs are introduced for the new hypotheses, The only additional mechanism required is the hypothesis weighting, which is applied in the schedule incongruity resolution step of the outer-loop. The following section of simulated results illustrates this through a more realistic decision-making scenario with multiple hypotheses.

5. Simulation results

This section contains a more nuanced application of JER scheduling sensor network actions to resolve multiple space situational awareness (SSA) hypotheses.

Table 4
Example Case 2: results based on true hypothesis realization.

Truth	Test 1				Test 2			
	Result	Pr _{pl} : A, B, C		$\tilde{H}_{JS}$	Result	Pr _{pl} : A, B, C		$\tilde{H}_{JS}$
A	1 - Pass	0.84,	0.08, 0.08	0.308	1 - Pass	0.98, 0.01, 0.01		0.055
B	1 - Fail	0.04,	0.48, 0.48	0.721	2 - Pass	0.02, 0.75, 0.02		0.389
C	1 - Fail	0.04,	0.48, 0.48	0.721	2 - Fail	0.07, 0.21, 0.71		0.468

5.1. Scenario motivation

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. Potential SSA needs include maintaining catalogs of space object state observations [34,35], detecting maneuvers or other anomalies [36], and estimating control modes or behavior [37,38]. Currently, there are over 20,000 trackable objects in the space object catalog [39,40] 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 [38]. The number of trackable space objects is continually growing with expanded use of small spacecraft technologies [41] 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 [42]. 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) [43]. 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” [38]. 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 [43,44].

Tracking techniques used in the space surveillance system still largely rely upon models and applications from the 1950s and 1960s [31], 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. 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 [45–47]. As the space object population increases, the amount of data required to maintain SSA also increases [48], which makes human-in-the-loop involvement in space surveillance particularly troublesome and motivates the development of autonomous sensor tasking capabilities.

SSA sensor tasking suffers from many competing objectives. For instance, maintaining a catalog of space object estimates requires observations of many different space objects. Information-maximizing methods,

as characterized through covariance estimates, minimize state estimate uncertainty for all catalog objects [8,34]. 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) [36,49], classification methods may employ taxonomies trained on representative space object feature sets to categorize space objects [50], 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 [37]. The competing objectives are generally not complementary, especially given limited sensor resources, so different objectives may prefer different tasking approaches.

5.2. Scenario description

Operators in a SSA decision-support environment receive notice from a space launch entity that a planned geostationary transfer orbit (GTO) insertion maneuver has experienced an anomaly. The anomaly is estimated to have occurred 5 min prior to the notification during a critical orbit-raising maneuver. In addition to existing priorities and tasking requirements, the two available sensors must be re-tasked to accommodate this new tasking priority. The objective is to re-acquire the space object and diagnose the anomaly to regain situation awareness.

This simulated scenario is a lower-dimensional (fewer sensors and objects) example of a very relevant research problem in SSA sensor tasking. 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 as the spacecraft was nominally bound for Geostationary Earth Orbit (GEO), a densely populated orbit regime with many high-value defense and telecommunications assets. The nominal transfer time from LEO to GEO is just over five hours, placing additional time-pressure on resolving the anomaly to complete conjunction analyses and alert other satellite operators. If the anomaly 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.

After a 5-minute delay between the anomaly event (at 2:00 UTC) and the beginning of the sensor tasking window (at 2:05 UTC), a simulated communication delay between the spacecraft operators and sensor network schedulers, the sensor operators have at most 10 minutes to gather observations and characterize the event. The sensor tasking time span is limited by observation constraints (e.g. short horizon-to-horizon 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 fully characterized by 02:15 UTC.

5.3. Dynamics

The nominal transfer orbit geometry is shown in Fig. 6. The primary spacecraft begins in a 1000 km altitude circular parking orbit. Space objects are propagated using Keplerian two-body dynamics to compute lines-of-sight to sensors. The sensor network is comprised of two 3-degree field-of-view electro-optical sensors, separated by 20 degrees in longitude for geometric diversity. Observations are taken with a




Fig. 6. Nominal GTO transfer orbit and target GEO orbit.

one-minute cadence and are simulated using a radiometric model, including simulated effects for background sky irradiance and atmospheric transmittance (e.g. cloud cover, atmospheric turbulence) [51] with illumination conditions estimated using a cannonball model.

The sensor-tasking time span is limited by observation constraints (e.g. short horizon-to-horizon times in LEO, eclipse, adverse weather), placing an upper time-limit on the hypothesis resolution. The sensors may change actions each minute, and a receding time-horizon of two minutes is used to evaluate potential action sequences.

5.4. Belief function models

A limited subset of potential failure modes is analyzed for illustrative purposes in this test case. As shown in Fig. 7, the anomaly is characterized at the subsystem level to determine root-cause. Since multiple point-of-failure events are exceedingly rare, an assumption is made that the anomaly results from a single point-of-failure, isolating the anomaly to one of these subsystems.

The hypotheses considered for this GTO insertion maneuver anomaly include: propulsion status, navigation status, and collision in LEO. To construct JER agent-pairs, each hypothesis is further decomposed into frames of discernment:

$$\Omega = \Omega_{man} \times \Omega_{prop} \times \Omega_{nav} \times \Omega_{coll} \quad (37)$$

Hypotheses are considered resolved if the normalized J-S entropy drops below the threshold value of $H_{JS,thr} = 0.05$.

The following sections describe each hypothesis, and the available evidence, in more detail.

5.4.1. Propulsion status

The propulsion status hypothesis, Ω_{prop} , yields a three-element frame:

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

Nominal propulsion status, $\omega_{prop,nom}$, represents the case where the propulsion subsystem is not 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.

5.4.2. Navigation status

The navigation status hypothesis, Ω_{nav} , yields a binary frame:

$$\Omega_{nav} = \{\omega_{nav,n}, \omega_{nav,a}\} \quad (38)$$

Nominal navigation status, $\omega_{nav,n}$, represents the case where the navigation subsystem is not the cause of the anomaly. Anomalous navigation, $\omega_{nav,a}$, results in an off-nominal transfer orbit due to pointing error, causing detection of the primary spacecraft off-track near the nominal GTO orbit.

5.4.3. Collision in LEO

The collision in LEO hypothesis, Ω_{coll} , yields 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) considered. For this illustrative example, three RSOs ($R = 3$) will be considered. 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 $j = 1, \dots, R$, results in debris in both orbits as well as missing nominal tracks for both object j and the primary spacecraft. Recall that explosion also generates debris in the LEO parking orbit, so the missing LEO object j and debris in its orbit differentiate the hypotheses. Nominal detection of an RSO refutes that RSO's collision proposition.

5.4.4. JER agent-pairs

The full problem considers each frame described in the decomposition above to investigate the cause of a maneuver anomaly. Each frame binary frame contributes one JER agent-pair, while each non-binary frame contributes $|\Omega|$ JER agent-pairs. Therefore, for this simulation, there are eight JER agent-pairs: three for propulsion status, one for navigation status, and four for collision in LEO.

5.5. Evidence to belief function mappings

Before utilizing any hypothesis-based approach, such as JER, analysts must develop a database of evidence-to-belief-function mappings that describe how gathered data affects hypothesis resolution. The process to develop these mappings is highly problem-specific, requiring the analyst to carefully consider what each potential successful or missed detection means with respect to each hypothesis. For instance, a missed detection of the nominal GTO orbit may indicate anomaly, but if the estimated electro-optical probability of detection [51] predicted a low chance of success, that evidence is vacuous and belief mass should be attributed to ignorance instead. In operation, each gathered piece of data is evaluated against the database to construct bpas that represent the effect of that piece of evidence on the hypotheses. The fusion of these evidence bpas with the a priori bpa generates an updated bpa representing an updated hypothesis knowledge.

In addition to explicit information carried by observational data, implicit knowledge about relationships between these frames can be imposed through conditional bpas [52]. While the propositions of a hypothesis are required by definition to be mutually exclusive, the hypotheses themselves do not need to be independent, and relationships between the hypotheses can be exploited through conditional bpas to resolve many hypotheses simultaneously. In this particular scenario, 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




Fig. 7. Possible causes for GTO insertion failure.




Fig. 8. Case 1: nominal maneuver (baseline).

maneuver status is likely anomalous. Conditional bpas are fused with the updated bpa, just like any other evidence bpa, to further update the hypothesis knowledge.

5.6. Case 1: nominal maneuver

As a baseline, the true proposition for this first test case is the nominal maneuver status. The resulting sensor tasking schedule is shown in Fig. 8(a) and (b) shows the normalized J-S entropy for each hypothesis. The resolutions of each proposition (belief and plausibility) are plotted in Fig. 9(a), (b), and (c).

The schedule in Fig. 8(a) indicates actions for each sensor at each time step, overlaid with target observability and tasking mode information. Together with Fig. 8(b), these figures provide an overview of the algorithm performance, allowing inspection of how each action taken (in the schedule) impacts the overall goal of resolving hypotheses. Note that, although the simulation begins at the anomaly epoch (2:00 UTC), the sensor-tasking window does not begin until five minutes later (2:05 UTC), as indicated on the schedule and results graphs. This is the same for all simulation cases, simulating a delay caused by required communication between the spacecraft operators and the sensor network operators.

For sensor 1, only the GTO target is valid for the first two steps. Its first attempted observation is missed, but the radiometric model for probability of detection confirms that the observation conditions (target near the horizon) contributed to this miss. As such, this missed observation does not affect the hypothesis resolution. Meanwhile, sensor 2 makes multiple successful observations early in the simulation,

including a task on the nominal GTO spacecraft that results in combined detection of the GTO object (supporting “nominal” propulsion and navigation statuses) and the collision objects “Coll 0” and “Coll 2”. These detections result in strong resolution of the respective collision propositions (see Fig. 9(c)) followed by successful detection of the “Coll 1” collision object (at 02:07). Sensor 2 also confirms the propulsion status “nominal” proposition by making a successful detection of the GTO object (at 02:05) and failing to detect the space object in its LEO parking orbit (at 02:06).

All hypotheses are resolved within the prescribed entropy tolerance within four steps (see Fig. 8(b)), so for the remainder of the simulation (from 02:09 onward) the sensors are free to perform other actions as necessary, as indicated by the none-of-the-above (NOTA) option and blue tasking mode. Using the sub-problem decomposition and efficient minimax search, JER only required a maximum of 271 sequence evaluations (including all agent-pairs and the combination schedule evaluations) in any iteration, far less than the theoretical brute-force maximum of 1024 evaluations.

5.7. Case 2: Propulsion non-start

In this 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 Fig. 10(a), and (b), respectively. The resolutions of each proposition are plotted in Fig. 11(a)–(c).

Similar to the baseline test case, strong evidence is available to confirm hypothesis resolutions quickly despite observability concerns. The




Fig. 9. Case 1: nominal maneuver (baseline), hypothesis resolutions (solid line for belief, dashed line for plausibility).




Fig. 10. Case 2: propulsion non-start.

initial action by sensor 2 misses an expected observation of the GTO object, and since observation conditions did not preclude this observation, this action contributes strong evidence to refute the “nominal” propulsion status proposition. Simultaneously, this action successfully detects an object in the LEO parking orbit, correlating it to the non-maneuvered trajectory and contributing strong evidence to confirm the “non-start” propulsion status proposition. Subsequent successful detections of each collision object (including a combined detection of “Coll 0” and “Coll 2” at 02:05) refute each collision proposition, and the navigation status is confirmed nominal by detection of the primary spacecraft in LEO.

Once again, all hypotheses are resolved within prescribed tolerances quickly, within four minutes, allowing the sensors to resume other tasking priorities as indicated by the “NOTA” option.

5.8. Case 3: propulsion system explosion

In this test case, a propulsion anomaly occurs resulting in an explosion, scattering debris in the LEO parking orbit. In total, five pieces of debris large enough for electro-optical detection are generated, originating from the LEO parking orbit at the anomaly epoch. The resulting sensor tasking schedule and hypothesis entropies are shown in Fig. 12(a)




Fig. 11. Case 2: propulsion non-start, hypothesis resolutions (solid line for belief, dashed line for plausibility).




Fig. 12. Case 3: propulsion explosion.

and (b), respectively. The resolutions of each proposition are plotted in Fig. 13(a)–(c).

This test case features weaker evidence, resulting in more actions required to reach adequate hypothesis resolution within prescribed tolerances. Initial attempts by both sensors to observe the object in GTO result in missed detections, contributing weak evidence toward both the anomalous propositions for both navigation status and propulsion status (non-start and explosion). The sensor network also initiates a search in GTO (as indicated by the red region on Fig. 12(a)) to confirm that

the navigation status is not anomalous, searching for the object in an off-nominal GTO state.

In the course of this search, several pieces of debris are detected, contributing evidence toward both the propulsive explosion and collision propositions. This initially inflates the belief in a collision with object “Coll 1” that is later refuted through positive detection of the “Coll 1” object in its nominal orbit. Evidence mounts toward the propulsive explosion proposition as further debris is detected in the LEO parking orbit, the target object is not found in GTO or LEO (refuting navigation anomaly and propulsion non-start), and each collision object is success-




Fig. 13. Case 3: propulsion explosion, hypothesis resolutions (solid line for belief, dashed line for plausibility).

fully detected (refuting any belief that the debris was generated by collision).

This test case serves as a prime example of the unbiased resolution focus of JER, as the prior induced by the initial debris detection is rejected by further evidence. Despite the weaker available evidence in this scenario, all hypotheses are resolved within tolerance by the end of the sensor tasking window.

5.9. Case 4: collision with object in LEO

In this test case, the true proposition is a collision in LEO with the object labeled “Coll 0.” This event generates multiple debris objects in both the LEO parking orbit and the nominal orbit of the collision object, with five unique detectable debris in each orbit. The resulting sensor tasking schedule and hypothesis entropies are shown in Fig. 14(a), and (b), respectively. The resolutions of each proposition are plotted in Fig. 15(a)–(c).

Similar to the propulsion explosion test case, detected debris increases belief in the collision hypothesis for object “Coll 0,” and successful detections of “Coll 1” and “Coll 2” refute those collision propositions. In this case, debris is detected throughout the simulation in both the LEO and “Coll 0” orbits, which differentiates the explosion and collision propositions, contributing weak evidence to refute the explosion proposition. A navigation anomaly is also ruled out through search of the GTO orbit, despite an initial slight belief in navigation anomaly, once again indicating an ability to overcome poor or incorrect prior knowledge.

As the simulation progresses, continued detections of debris objects builds to a strong resolution of collision with object “Coll 0” and nearly-refuted propulsion explosion. At the end of the tasking window, the propulsion status hypothesis is not fully resolved within the prescribed tolerance, but the collision status resolution paints a clear picture of debris in both the LEO parking and collision object orbits.

5.10. Comparison to entropy-greedy scheduler

Comparing this SSA application of JER to other sensor-tasking approaches is difficult, primarily because SSA sensor tasking algorithm development currently focuses on the reduction of space object state

estimate uncertainty. Characterized through state covariance estimates, the state uncertainty is a natural scheduling algorithm metric due to its direct relationship to some decision-making questions (i.e. collision). However, not all decision-making hypotheses can be readily transposed to state covariance thresholds; for instance, the active mode of a space object is inferred from light-curve inversion [37], not state estimate uncertainty. This makes comparisons to JER in SSA hypothesis resolution tasks difficult, as the direct emphasis on varied decision-maker objectives is unique to SSA sensor tasking.

Instead, a brute-force entropy-greedy scheduler was implemented for comparison to the JER approach. The entropy-greedy scheduler evaluates all valid action sequences over the scheduler horizon and selects the action sequence that minimizes the weighted-sum entropy. This represents the brute-force solution to the hypothesis-based evidence-gathering optimization problem in Eq. (15), presented to analyze proposed computational complexity and bias-related improvements of the JER approach.

The biggest immediate difference between the approaches is the number of sequence evaluations required. Even in these low-dimensional scenarios, the brute-force evaluation of all possible tasking sequences (two sensors, five targets, two-step horizon) requires 1,024 sequence evaluations each iteration for all test cases. In comparison, for the nominal maneuver scenario (Case 1), recall that JER only required a maximum of 271 sequence evaluations (including all agent-pairs and the combination schedule evaluations) in any iteration.

While the entropy-greedy scheduler sometimes performs identically to JER in hypothesis resolution (as in Case 1 and 2), it predictably struggles with confirmation bias. The comparison scenario presented is identical to the propulsion anomaly in Case 3. The resulting sensor tasking schedule and hypothesis entropies are shown in Fig. 16(a), and (b), respectively. The resolutions of each proposition are plotted in Figs. 13(a), 17(b) and (c).

As with Case 3 above, detection of debris early in the simulation contributes evidence toward both the (correct) explosion proposition and the (incorrect) collision propositions. With JER, the alternating-agent scheme overcomes the incorrect collision prior by searching for confirming evidence for each collision proposition. However, in the entropy-greedy approach, once the other propositions are ruled out and only the




Fig. 14. Case 4: collision in LEO.




Fig. 15. Case 4: collision in LEO, hypothesis resolutions (solid line for belief, dashed line for plausibility).

“Coll 1” proposition remains, the algorithm is satisfied with its hypothesis resolution. This results in an incorrect resolution of the collision status hypothesis.

This is a phenomenon also experienced in previous work using DST for sensor tasking [53], where the mutually exclusive and collectively exhaustive nature of the propositions does not always encourage positive confirmation of the hypothesis resolution. The entropy-greedy results, contrasted with the JER results, further underscore the impact of the alternating-agent scheme in rejecting confirmation bias.

5.11. 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 computation-




Fig. 16. Entropy-greedy scheduler: propulsion explosion.




Fig. 17. Entropy-greedy scheduler: propulsion explosion, hypothesis resolutions (solid line for belief, dashed line for plausibility).

ally constrained in a brute-force approach, even for this relatively low-dimensional example.

The test cases range from clear scenarios with strong evidence to anomalous scenarios with weak and ambiguous evidence. The nominal maneuver scenario (Case 1) provides a baseline for comparison, quickly resolving the hypotheses with strong evidence. The propulsion non-start scenario (Case 2) shows an ability to ingest both weak evidence (missed detections from GTO) and strong evidence (successful detections in LEO) to explore the hypotheses efficiently. The propulsion explosion and collision scenarios (Cases 3 and 4) highlight the ability to avoid confirma-

tion bias induced by poor prior knowledge or ambiguous evidence by continuing to seek evidence to reject the incorrect propositions.

In comparison, the entropy-greedy approach struggled to overcome confirmation bias, resulting in an incorrect hypothesis resolution for the propulsion explosion test case shown. The incorrect resolution is avoided by JER through the alternating-turn adversarial approach of the agent-pairs, accounting for and avoiding confirmation bias by ensuring an advocate exists for each proposition. Despite the limited nature of this SSA application scenario (relatively few hypotheses, sensors, and available actions), the JER approach demonstrates a marked

difference over an entropy-greedy approach, indicating that the intuition to gather the estimated strongest-available evidence can lead to incorrect resolutions due to confirmation bias. These results provide a more thorough and realistic JER application illustration than the medical examples, and future work will continue to refine this application to improve realism through increased numbers and difficulties of hypotheses and more sensors with more available actions. In addition to addressing problems with many more hypotheses and actions, a diverse sensor network presents additional challenges, including asynchronous operations that must be balanced in the agent-pair implementation.

Further comparisons to existing and proposed SSA sensor tasking methods are confounded by the lack of clear translations from decision-maker hypotheses to state estimate covariance, the commonly-used SSA sensor tasking optimization metric. JER can be applied to any decision-making hypothesis (not just those related to covariance), and this difference in optimization focus (hypothesis-resolution vs. state estimate uncertainty) make JER too distinct for direct performance comparisons. However, comparisons between the operator decision-making and cognitive effects of covariance- and hypothesis-based methods are relevant to the development of proposed sensor tasking approaches. Concurrent work has investigated cognitive support, situation awareness, workload, and performance effects of the two scheduling approaches [54], but since this paper focuses on the technical details of the development of the JER algorithm, the cognitive effects comparisons are considered out of scope of this development.

6. Conclusion

The proposed Judicial Evidential Reasoning (JER) evidence-gathering framework arranges decision-maker questions as rigorously testable hypotheses to enable predictive evidence-gathering for hypothesis resolution. 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 the principle of equal effort through 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 evidence-gathering 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 provided example cases illustrate the application of both the JER agent-pairs and the overall JER schedule manager approach to evidence-gathering. 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. The JER approach also compares well against an entropy-greedy approach that focuses actions on the most-probable propositions only, avoiding improper hypothesis resolution caused by confirmation bias.

Continuing research investigates the human cognitive effects of a JER-like hypothesis-based evidence gathering approach to further develop decision support systems that effectively support human-in-the-loop decision-making [54]. Additionally, the principle of equal effort is currently based on a heuristic in lieu of well-developed metrics for confirmation bias. Further research in defining methods for measuring confirmation bias would provide a more quantitative means for estimating and mitigating confirmation bias. A related concern is the consid-

eration of asynchronous evidence-gathering tasks. If evidence-gathering sources in a given network operate at different rates, care must be taken to ensure that one faster sensor's evidence does not overwhelm evidence gathered by other slower sensors. Further comparison studies can also be conducted using different measures of bpa entropy to examine different approaches to quantifying and reducing hypothesis uncertainty.

The developed JER algorithm allows the application of a hypothesis abstraction, quantified ambiguity, and considerations for confirmation bias in order to rigorously address decision-maker hypotheses through the timely application of appropriate evidence. Through these combined emphases, JER enables predictive evidence-gathering for hypothesis resolution.

References

- [1] M.R. Endsley, Situation awareness global assessment technique (SAGAT), in: IEEE 1988 National Aerospace and Electronics Conference, 1988, doi:10.1109/NAECON.1988.195097.
- [2] 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, 2755, 1996, pp. 304–312, doi:10.1117/12.243172.
- [3] D. Shen, G. Chen, K. Pham, E. Blasch, A trust-based sensor allocation algorithm in cooperative space search problems, Sensors and Systems for Space Applications IV, Vol. 8804, Orlando, FL, 2011, doi:10.1117/12.882904.
- [4] J. Liu, M. Chu, J. Reich, Multitarget tracking in distributed sensor networks, IEEE Signal Process. Mag. 24 (3) (2007) 36–46, doi:10.1109/MSP.2007.361600.
- [5] S.S. Blackman, Multiple hypothesis tracking for multiple target tracking, IEEE Aerosp. Electr. Syst. Mag. 19 (1) (2004) 5–18, doi:10.1109/MAES.2004.1263228.
- [6] S. Gehly, B. Jones, P. Axelrad, Sensor allocation for tracking geosynchronous space objects, J. Guid., Control, Dyn. 41 (1) (2018) 149–163, doi:10.2514/1.G000421.
- [7] C. Frueh, H. Fielder, J. Herzog, Heuristic and optimized sensor tasking observation strategies with exemplification for geosynchronous objects, J. Guid., Control, Dyn. 41 (5) (2018) 1036–1048, doi:10.2514/1.G003123.
- [8] I.I. Hussein, K.J. DeMars, C. Frueh, 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, 2012.
- [9] A.D. Jaunzemis, M.J. Holzinger, Evidential reasoning applied to single-object loss-of-custody scenarios for telescope tasking, in: 26th AAS/AIAA Spaceflight Mechanics Conference, 2016.
- [10] A.D. Jaunzemis, D. Minotra, M.J. Holzinger, K.M. Feigh, M.W. Chan, P.P. Shenoy, Judicial evidential reasoning for decision support applied to orbit insertion failure, in: 1st IAA Conference on Space Situational Awareness, Orlando, FL, 2017.
- [11] G. Shafer, A Mathematical Theory of Evidence, Princeton University Press, 1976.
- [12] G. Shafer, R. Logan, Implementing Dempster's rule for hierarchical evidence, Artif. Intell. 33 (1987) 271–298.
- [13] A.P. Dempster, The Dempster–Shafer calculus for statisticians, Int. J. Approximate Reasoning 48 (2008) 365–377.
- [14] T. Denoeux, A k-nearest neighbor classification rule based on Dempster–Shafer theory, IEEE Trans. Syst., Man, Cybern. 25 (5) (1995) 804–813, doi:10.1109/21.376493.
- [15] S.L. Hegarat-Masle, I. Blouch, D. Vidal-Madjar, Application of Dempster–Shafer evidence theory to unsupervised classification in multisource remote sensing, IEEE Trans. Geosci. Remote Sens. 35 (4) (1997) 1018–1031, doi:10.1109/36.602544.
- [16] C.R. Parikh, M.J. Pont, N.B. Jones, Application of Dempster–Shafer theory in condition monitoring applications: a case study, Pattern Recognit. Lett. 22 (6–7) (2001) 777–785, doi:10.1016/S0167-8655(01)00014-9.
- [17] B.-S. Yang, K.J. Kim, Application of Dempster–Shafer theory in fault diagnosis of induction motors using vibration and current signals, Mech. Syst. Signal Process. 20 (2) (2006) 403–420.
- [18] W.F. Caselton, W. Luo, Decision making with imprecise probabilities: Dempster–Shafer theory and application, Water Resour. Res. 28 (12) (1992) 3071–3083.
- [19] D. Frey, M. Rosch, Information seeking after decisions, Pers. Soc. Psychol. Bull. 10 (1) (1984) 91–98, doi:10.1177/0146167284101010.
- [20] E. Jonas, S. Schulz-Hardt, D. Frey, N. Thelen, Confirmation bias in sequential information search after preliminary decisions: an expansion of dissonance theoretical research on selective exposure to information, J. Pers. Soc. Psychol. 80 (4) (2001) 557–571.
- [21] M.B. Cook, H.S. Smallman, Human factors of the confirmation bias in intelligence analysis: decision support from graphical evidence landscapes, Hum. Factors 50 (5) (2008) 745–754, doi:10.1518/001872008X354183.
- [22] D. Ellsberg, Risk, ambiguity, and the Savage axioms, Q. J. Econ. 75 (4) (1961) 643–669.
- [23] B.R. Cobb, P.P. Shenoy, On the plausibility transformation method for translating belief function models to probability models, Int. J. Approximate Reasoning 41 (3) (2006) 314–330.
- [24] A.P. Dempster, Upper and lower probabilities induced by a multivalued mapping, Ann. Math. Stat. 4 (1) (1967) 325–339.
- [25] R. Jirousek, P.P. Shenoy, A new definition of entropy of belief functions in the Dempster–Shafer theory, Int. J. Approximate Reasoning 92 (1) (2018) 49–65.
- [26] P. Smets, R. Kennes, The transferable belief model, Artif. Intell. 66 (2) (1994) 191–234.

- [27] G. Shafer, The combination of evidence, *Int. J. Intell. Syst.* 1 (1986) 155–179.
- [28] 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, 2016.
- [29] R.R. Yager, Arithmetic and other operations on Dempster–Shafer structures, *Int. J. Man Mach. Stud.* 25 (4) (1986) 357–366, doi:10.1016/S0020-7373(86)80066-9.
- [30] Y. Deng, Deng entropy, *Chaos, Solitons and Fractals* 91 (2016) 549–553.
- [31] F.R. Hoofs, P.W. Schumacher, History of analytical orbit modeling in the U.S. space surveillance system, *J. Guid., Control, Dyn.* 27 (2004) 174–185, doi:10.2514/1.9161.
- [32] P.P. Shenoy, G. Shafer, Axioms for probability and belief-function propagation, in: R.D. Shachter, T. Levitt, J.F. Lemmer, L.N. Kanal (Eds.), *Uncertainty in Artificial Intelligence*, 4, Elsevier, Amsterdam, 1990, pp. 169–198, doi:10.1016/B978-0-444-88650-7.50019-6.
- [33] D.E. Knuth, R.W. Moore, An analysis of alpha-beta pruning, *Artif. Intell.* 6 (4) (1975) 293–326, doi:10.1016/0004-3702(75)90019-3.
- [34] K.J. DeMars, I.I. Hussein, C. Frueh, M.K. Jah, R. Scott Erwin, Multiple-object space surveillance tracking using finite-set statistics, *J. Guid., Control, Dyn.* (2015) 1–16, doi:10.2514/1.G000987.
- [35] T.A. Hobson, *Sensor Management for Enhanced Catalogue Maintenance of Resident Space Objects*, University of Queensland Australia, 2014 Ph.D. thesis.
- [36] A.D. Jaunzemis, M.V. Mathew, M.J. Holzinger, Control cost and Mahalanobis distance binary hypothesis testing for spacecraft maneuver detection, *J. Guid., Control, Dyn.* 39 (9) (2016) 2058–2072, doi:10.2514/1.G001616.
- [37] 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, *J. Guid. Control Dynam.* 41 (1) (2018) 164–170, doi:10.2514/1.G002223.
- [38] 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, *NATO Symposium on Considerations for Space and Space-Enabled Capabilities in NATO Coalition Operations*, 2016, doi:10.13140/RG.2.1.3057.1128.
- [39] 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.
- [40] Orbital Debris Quarterly News, Technical Report, NASA Orbital Debris Program Office, 2017.
- [41] 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)*, 2015, doi:10.1109/RAST.2015.7208459.
- [42] T. Kelso, Analysis of the Iridium 33 Cosmos 2251 collision, in: *Advanced Maui Optical and Space Surveillance Technologies*, Maui, HI, 2009.
- [43] J.D. Ianni, D.L. Alea, 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.
- [44] P.A. Brown, Promoting the safe and responsible use of space: toward a 21st century transparency framework, *High Fron.* 5 (1) (2008).
- [45] 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, *Astron. J.* 147 (2014) 111–124, doi:10.1088/0004-6256/147/5/111.
- [46] T. Hardy, S. Cain, J. Jeon, T. Blake, Improving space domain awareness through unequal-cost multiple hypothesis testing in the space surveillance telescope, *Appl. Opt.* 54 (17) (2015) 5481–5494, doi:10.1364/AO.54.005481.
- [47] T.S. Murphy, M.J. Holzinger, B. Flewelling, Space object detection in images using matched filter bank and Bayesian update, *J. Guid., Control, Dyn.* 40 (3) (2017) 497–509, doi:10.2514/1.G001934.
- [48] T. Blake, M. Sanchez, J. Krassner, M. Georgen, S. Sundbeck, *Space Domain Awareness*, Technical Report, DARPA, 2012.
- [49] M.J. Holzinger, D.J. Scheeres, K.T. Alfriend, Object correlation, maneuver detection, and characterization using control distance metric, *J. Guid. Navigation Control* 35 (4) (2012) 1312–1325, doi:10.2514/1.53245.
- [50] 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, 2013.
- [51] R.D. Coder, M.J. Holzinger, Multi-objective design of optical systems for space situational awareness, *Acta Astronaut.* (2016), doi:10.1016/j.actaastro.2016.07.008.
- [52] G. Shafer, Belief functions and parametric models, *J. R. Stat. Soc. Ser. B* 44 (3) (1982) 322–352.
- [53] A.D. Jaunzemis, M.J. Holzinger, K.K. Luu, Sensor tasking for spacecraft custody maintenance and anomaly detection using evidential reasoning, *J. Aerosp. Inf. Syst.* (2018), doi:10.2514/1.1010584.
- [54] A.D. Jaunzemis, D. Minotra, K.M. Feigh, M.J. Holzinger, Cognitive systems engineering applied to decision support in ssa, in: *19th Advanced Maui Optical and Space Surveillance Technologies Conference (Accepted)*, Kihei, HI, 2018.