



FUSION 2018

21st International Conference on Information Fusion

10-13 JULY 2018 | UNIVERSITY OF CAMBRIDGE | UK

[2018 21st International Conference on Information Fusion \(FUSION\)](#) took place
10-13 July 2018 in Cambridge, United Kingdom

IEEE catalog number: CFP18FUS-USB

ISBN: 978-0-9964527-7-9

Copyright and Reprint Permission: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limit of U.S. copyright law for private use of patrons those articles in this volume that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923. For other copying, reprint or republication permission, write to IEEE Copyrights Manager, IEEE Operations Center, 445 Hoes Lane, Piscataway, NJ 08854. All rights reserved. Copyright © 2018 by IEEE.

Evidence Gathering for Hypothesis Resolution using Judicial Evidential Reasoning

Andris D. Jaunzemis
and Marcus J. Holzinger
Georgia Institute of Technology
Atlanta, Georgia 30332
adjaunzemis@gatech.edu
holzinger@ae.gatech.edu

Moses W. Chan
Lockheed Martin Corporation
Sunnyvale, CA, USA, 94089
moses.w.chan@lmco.com

Prakash P. Shenoy
University of Kansas
Lawrence, KS, USA,
pshenoy@ku.edu

Abstract—Realistic decision-making often occurs with insufficient time to gather all possible evidence before a decision must be rendered, requiring efficient processes for prioritizing between candidate action sequences. The proposed Judicial Evidential Reasoning framework encodes decision-maker questions as rigorously testable hypotheses and proposes actions to resolve the hypotheses in the face of ambiguous, incomplete, and uncertain evidence. Dempster-Shafer theory is applied to model hypothesis knowledge and quantify ambiguity, and an equal-effort heuristic is proposed time-efficiency and impartiality to combat confirmation bias. This work includes derivation of the generalized formulation, computational tractability considerations for improved performance, several illustrative examples, and sample 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.

I. 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”. The evidence gathering or scheduling problem addresses how to obtain, process, and utilize evidence to improve understanding of the state of the environment [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).

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. This motivates the use of evidential reasoning approaches, such as Dempster-Shafer theory, to quantify ambiguity and realistically model decision-making processes.

Another concern in evidence-gathering is confirmation bias, a preferential tendency to gather evidence that confirms 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.

This work develops the Judicial Evidential Reasoning (JER) approach, improving upon preliminary work [3] by generalizing the framework to any hypothesis-resolution domain while also improving computational tractability. 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 sensor tasking problem using adversarial optimization techniques. Several examples and simulations are provided to illustrate the application of this framework.

II. BACKGROUND

This section introduces background material relevant to the theoretical developments.

A. Ambiguity Aversion

Bayesian probability theory is a prevailing methodology for reasoning in uncertain domains by modeling knowledge about propositions using true-or-false probabilities [4]. However, probability theory struggles to express evidence ambiguity due to noise, poor observability, or ignorance on the part of the evidence source.

A phenomenon known as *ambiguity aversion* shows that decision-makers overwhelmingly prefer known risks to unknown risks, making ambiguity a major concern in modeling knowledge states [5]. Ellsberg’s paradox shows that humans overwhelmingly prefer to bet on an equal-probability prior than a fully ambiguous prior. In Bayesian probability, both priors would be represented as the same (equal-probability) probability mass function, so information on the presence or lack of ambiguity is not adequately conveyed to the decision-maker. This deficiency in Bayesian probability theory that has a significant impact in human decision-making contexts and motivates alternative methodologies such as evidential reasoning [4].

B. Dempster-Shafer Theory

Dempster-Shafer theory is considered more expressive than probability theory in representing ambiguity or ignorance [6].

This is accomplished by allowing assignment of belief to non-singleton propositions, admitting ambiguity on the part of the expert when necessary.

In classical Dempster-Shafer theory, mutually exclusive and collectively exhaustive propositions form the frame of discernment, Ω , for a given hypothesis. Elements of the powerset 2^Ω , the set of all subsets of Ω , are referred to as propositions.

1) *Belief Functions*: A basic probability assignment (bpa) m , as defined in Eqn. (1), maps a belief mass to each proposition:

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

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

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

The constraint in Eqn. (2) enforces the mutually exclusive and collectively exhaustive properties, while the constraint in Eqn. (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 [7]. Various bpa combination rules have been developed to fuse evidence from multiple sources into one bpa [8], with the most common being Dempster's conjunctive rule [7].

2) *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 [4]. Multiple methods exist for translating between Dempster-Shafer belief functions and probability models, allowing the use of Bayesian decision theory. One such method, the plausibility probability transformation defined in Eqn (4), is more consistent with Dempster's rule [4]:

$$\text{Pr}_{\text{plm}}(x) = K^{-1} \text{Pl}_m(\{x\}), \quad (4)$$

$$\text{where } K = \sum_{x \in \Omega} \text{Pl}_m(\{x\}) \quad (5)$$

Entropy is also frequently used as an information content measure in both probabilistic and evidential reasoning. Multiple proposed definitions of entropy for Dempster-Shafer theory are summarized in [6]. The Jirousek-Shenoy (J-S) definition of entropy combines Shannon and Dubois-Prade entropy to capture both conflict and non-specificity, appropriately modeling hypothesis knowledge to handle ambiguity aversion [6]. Maximum J-S entropy is attained by a vacuous bpa, which is the bpa where both conflict and ambiguity are highest as all belief mass is assigned to the frame: $m(\Omega) = 1$. A decision-maker wants to minimize conflict and non-specificity in hypothesis resolution, so minimizing J-S entropy is a reliable metric for strong hypothesis resolution considering ambiguity aversion.

III. GENERALIZED JER APPROACH

The following section further develops the Judicial Evidential Reasoning (JER) approach to evidence-gathering hy-

pothesis resolution. This work generalizes the framework first presented in [3] to arbitrary hypothesis-resolution domains and addresses several computational tractability considerations. The JER approach hinges upon three primary considerations: hypothesis abstraction, ambiguity aversion, and unbiased hypothesis resolution.

A. Hypothesis Abstraction

Many evidence-gathering approaches (e.g. sensor network tasking) operate on maintaining a low overall state uncertainty; however, it may not be readily apparent to a decision-maker how reducing state uncertainty 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.

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 [9]. Forcing an operator to switch between different levels of the abstraction, approaching the problem at multiple different levels of detail, leads to increased frustration and workload and decreased situation awareness. Designing a decision-support system that directly conveys expected hypothesis resolution from candidate actions allows the human decision-maker to focus effort on strategic cognitive tasks.

B. General 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, 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.

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 at time t_k are described through the Cartesian product:

$$\mathbb{A}_k = \mathbb{A}_{1,k} \times \dots \times \mathbb{A}_{m,k} \quad (6)$$

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):

$$\mathcal{A}_{s,1:T} = (A_{s,1}, \dots, A_{s,T}), \quad s = 1, \dots, m \quad (7)$$

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

$$\mathcal{A}_{1:T} = \{\mathcal{A}_{1,1:T}, \dots, \mathcal{A}_{m,1:T}\} \quad (8)$$

This set contains an action sequence for each of the m actors and thus fully defines all actions 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:

$$\mathbb{A}_{1:T} = \mathcal{A}_{1:T} \times \dots \times \mathcal{A}_{1:T} \quad (9)$$

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:

$$J_T : (\Omega, \mathcal{W}; \mathcal{A}_{1:T}) \mapsto \mathbb{R} \quad (10)$$

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.

The general hypothesis-based evidence-gathering optimization problem is:

$$\mathcal{A}_{1:T}^* = \arg \min_{\mathcal{A}_{1:T} \in \mathbb{A}_{1:T}} J_T(\Omega, \mathcal{W}; \mathcal{A}_{1:T}) \quad (11)$$

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 $A_{s,\cdot} \in \mathcal{A}_{s,1:T}$ 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.

C. Hypothesis Resolution and Entropy

Hypothesis resolution refers to the goal of determining which proposition is true from the set of propositions in the frame of discernment. Recall that J-S entropy [6] quantifies both conflict and non-specificity in hypothesis knowledge, providing an apt minimization objective for strong hypothesis resolution.

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 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:

$$m_{i,j,k} : 2^{\Omega_i} \mapsto [0, 1] \quad (12)$$

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

$$\tilde{m}_{i,k} = \bigoplus_{j=1}^p m_{i,j,k} \quad (13)$$

Recall that Dempster's rule is associative and commutative, meaning the combination can be done sequentially in any order [10]. However, Dempster's rule is not idempotent, so the pieces of evidence must be independent to avoid artificially inflating 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.

The resultant knowledge state for hypothesis Ω_i , incorporating all evidence from time from t_0 to t_k is denoted as:

$$m_{i,k}^+ = m_{i,k}^- \oplus \tilde{m}_{i,k} \quad (14)$$

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 .

1) *Normalized Jirousek-Shenoy Entropy*: The resolution of hypothesis Ω_i based on bpa m_i , as measured through J-S entropy [6], is defined as:

$$H_{JS}(m_i) = \sum_{x \in \Omega_i} \text{Pl}_{P_{m_i}}(x) \log_2 \left(\frac{1}{\text{Pl}_{P_{m_i}}} \right) + \sum_{A \in 2^{\Omega_i}} m_i(A) \log_2(|A|) \quad (15)$$

where the first summation term, related to Shannon entropy, quantifies conflict and the second summation term, called Dubois-Prade entropy, quantifies non-specificity.

To use entropy as a cost function while accounting for hypotheses with different numbers of propositions, the normalized Jirousek-Shenoy entropy is defined as follows:

$$\tilde{H}_{JS}(m_i) = \frac{H_{JS}(m_i)}{2 \log_2(|\Omega_i|)} \quad (16)$$

where m_i is the bpa representing hypothesis Ω_i .

2) *Optimization Formulation*: To accomplish the goal of minimizing hypothesis conflict and non-specificity, the normalized entropy defined in Eqn. (16) is employed as the cost function for the optimization problem in Eqn. (11).

$$\mathcal{A}_{1:T}^* = \arg \min_{\mathcal{A}_{1:T} \in \mathbb{A}_{1:T}} \sum_{i=1}^{|\Omega|} w_i \tilde{H}_{JS}(\hat{m}_{i,T}) \quad (17)$$

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} . The 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) *Computational Complexity*: The general formulation in Eqn. (17) suffers from a number of practical issues in implementation. Most notably, computational complexity scales with the number of hypotheses $|\Omega|$, the number of actors m , the number of valid actions for each actor n_m , and the time horizon T :

$$\mathcal{O} \left(\prod_{t=1}^T \left(\prod_{s=1}^m n_{s,t} \right) \right) \quad (18)$$

where $n_{s,t}$ is the number of valid actions for the s^{th} actor at time t_{k+t} . Using the worst-case number of valid actions $n = \max(n_{s,t} \mid s = 1, \dots, m; t = 1, \dots, T)$:

$$\mathcal{O}\left((n^m)^T\right) \quad (19)$$

provides an upper bound on the brute-force complexity.

D. Implementation Considerations

This section modifies the general optimization approach in Eqn. (17) to arrive at a computationally-tractable solution by decomposing the problem into independent sub-problems. An additional concern with the entropy-reduction algorithm is evidence-gathering bias induced by prior information and ambiguity. Adversarial optimization is applied to reduce action sequence evaluations and combat confirmation bias.

1) *Unbiased Hypothesis Resolution*: Confirmation bias is a cognitive phenomenon where prior belief causes fixation on a particular proposition. If biased, the human may favor evidence that confirms prior beliefs and overlook conflicting evidence [11] or interpret ambiguous evidence in favor of prior beliefs. Similar to human cognitive fixation, socio-technical systems can also exhibit confirmation bias. Prior information has the potential to skew future evidence-gathering actions, inducing fixation through measurement noise, sensor bias, or other sources of uncertainty.

Avoiding fixation and confirmation bias adds a competing objective to the requirement of minimizing hypothesis entropy. Quantifying confirmation bias is an active area of research, with cognitive sciences researchers using various measures comparing selection of supporting versus refuting evidence [11], [12]. One such measurement is the difference between numbers of selected supporting and refuting evidence elements [13], 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.

2) *Sub-Problem Definition*: The primary intuition that allows decomposition of the entropy-reduction approach in Eqn. (17) is that not all actions contribute evidence related to all hypotheses. The sub-problems can be solved 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 $\mathbb{A}_{s,k,i}$:

$$\mathbb{A}_{s,k,i} \subseteq \mathbb{A}_{s,k} \quad (20)$$

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

$$\mathbb{A}_{1:T,i} \subseteq \mathbb{A}_{1:T} \quad (21)$$

The sub-problem optimization objective is first represented using a generic cost function:

$$J_{T,i} : (\Omega_i; \mathcal{A}_{1:T,i}) \mapsto \mathbb{R} \quad (22)$$

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

$$\mathcal{A}_{1:T,i}^* = \arg \min_{\mathcal{A}_{1:T,i} \in \mathbb{A}_{1:T,i}} J_{T,i}(\Omega_i; \mathcal{A}_{1:T,i}) \quad (23)$$

This sub-problem decomposition approach allows for parallel computation of action sequence sets for each agent-pair.

3) *Combating Confirmation Bias*: The adversarial optimization minimax technique is employed to reduce confirmation bias. In minimax optimization, 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.

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. 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:

$$\mathcal{A}_{1:T,i}^* = \arg \max_{\mathcal{A}_{1:T,i} \in \mathbb{A}_{1:T,i}} \min \Pr_{\text{pl}}(\theta; m_i | \mathcal{A}_{1:T,i}) \quad (24)$$

where $m_i | \mathcal{A}_{1:T,i}$ is the estimated bpa for hypothesis Ω_i resulting from the proposed action sequence $\mathcal{A}_{1:T,i}$. 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 $\{-\theta\}$ has full belief. When the critic agent is active, its goal is to maximize belief in the alternative proposition ($\{-\theta\}$) or equivalently minimize belief in the null proposition ($\{\theta\}$). Therefore, the formulation simply flips to a minimax optimization:

$$\mathcal{A}_{1:T,i}^* = \arg \min_{\mathcal{A}_{1:T,i} \in \mathbb{A}_{1:T,i}} \max \Pr_{\text{pl}}(\theta; m_i | \mathcal{A}_{1:T,i}) \quad (25)$$

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

4) *Resolving Combined Schedule Incongruity*: It is possible, or even likely, that the sub-problem optimal schedules for two or more agent-pairs use the same actor for different actions. These incongruities are resolved by choosing the actions that lead to the strongest hypothesis resolution using entropy.

Using the set of actions from all sub-problem optimal sequences $\mathbb{A}_{1:T}^*$, all possible combinations form candidate congruous action sequences. The evaluation criterion for selecting the optimal combined schedule is the weighted-sum of entropy:

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

where w_i is the weighting for the i^{th} hypothesis, and $\tilde{H}_{JS}$ is the normalized J-S entropy as defined in Eqn. (16). Since Jirousek-Shenoy 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.

5) *Efficient Minimax Optimization*: To further reduce the number of action sequences evaluated, the alternating-agent minimax formulation of the sub-problems can be further exploited using alpha-beta pruning [14]. In naive minimax (or maximin) optimization, the number of evaluations grows exponentially with the number of valid actions and the search depth, as in Eqns. (18) and (19). However, some action sequences can be pruned to quickly eliminate costly or infeasible options.

Pruning the known sub-optimal branches reduces the number of required sequence evaluations while still arriving at the same optimal solution as naive minimax. In an ideal case, the computational complexity reduces to Eqn. (27), a significant improvement over the brute-force complexity in Eqn. (19).

$$\mathcal{O}\left(\sqrt{(n^m)^T}\right) \quad (27)$$

While this idealized complexity may not be fully realized in application, alpha-beta pruning is likely to eliminate unnecessary searches to provide a more efficient minimax search.

6) *Hypothesis Pruning via Entropy Stopping Condition*: A final computational consideration is the pruning of resolved hypotheses. Once sufficient evidence has been gathered to resolve a hypothesis, it is beneficial to remove that hypothesis from consideration for future tasking evaluations. Decision-makers should be able to indicate an acceptable level of conflict and ambiguity, manifesting as J-S entropy thresholds $\tilde{H}_{th}(m_i)$ for each hypothesis Ω_i . If the entropy for a given

hypothesis falls below this threshold, that hypothesis is considered adequately resolved and action sequences related to that hypothesis no longer need to be considered.

IV. EXAMPLES

This section contains illustrative examples of the JER approach using simplified medical diagnosis situations with constraints enforced such that not all actions may be taken within the diagnosis window (i.e. time-critical decision-making).

A. Case 1: Single JER Agent-Pair

The first example involves a single hypothesis with two 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. Three tests are available to inform this diagnosis, but time and cost constraints limit the number of tests to two. Therefore, the goal is to determine which two tests result in a strong-but-unbiased resolution. Table I lists basic probability assignments (bpas) for each available test, functions of known statistics on the test such as false alarm rate.

TABLE I
EXAMPLE CASE 1: BASIC PROBABILITY ASSIGNMENTS FOR 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

Fig. 1 shows the tree of all possible test sequences, easily visualized due to the low dimensionality of this example. Traversing down the tree, each successive level alternates the active agent between advocate and critic (of A) for a two-step minimax optimization on $\Pr_{pl}(A)$. Each edge of this tree is a candidate action (test), and each terminating node denotes the plausibility probability of A as a result of the two actions leading to it. The non-terminating nodes display the chosen node from below based on the active agent at that step (min or max).




Fig. 1. Example Case 1: Evaluation tree with alpha-beta pruning

The minimax-optimal test sequence (tests 1 followed by test 2) is highlighted in blue, resulting in $\Pr_{pl}(A) = \Pr_{pl}(\neg A) =$

0.5. 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. The nodes marked with an X in Fig. 1 are pruned before evaluation via alpha-beta pruning, reducing the required evaluations from nine to seven.

B. Case 2: Multiple JER Agent-Pairs

The second example involves a single hypothesis with three propositions, yielding the following frame of discernment: $\{\theta_1, \theta_2, \theta_3\}$. For ease of notation, define the corresponding propositions as: $A = \{\theta_1\}$, $B = \{\theta_2\}$, $C = \{\theta_3\}$. Since the frame of discernment is non-binary, three JER agent-pairs are used. 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. Three tests are available to inform this diagnosis, but time and cost constraints limit the number of tests to two. Table II lists basic probability assignments (bpas) for each test, with two possible outcomes: pass and fail. 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 $B \cup C = \neg A$.

TABLE II

EXAMPLE CASE 2: BASIC PROBABILITY ASSIGNMENTS FOR 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

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 pass while tests 2 and 3 will fail, contributing belief to proposition A .

After the sub-problem test schedules are optimized, the estimated bpa result is computed for each combination sequence, and the test sequence resulting in the lowest entropy is selected. The first test action implemented for that iteration, and the agent-pairs that requested that action (i.e. the active agent-pairs) are flipped so that, in the next iteration, the critic agent is active first.

In test execution, the test bpa is determined by the true realization of the hypothesis. The prior bpa is combined with the test bpa after each iteration, and the procedure above is repeated using this updated bpa as the prior.

Table III shows the resulting test sequence and probabilities (using the plausibility transformation) after each iteration. Three different cases are run, each with a different true realization of the hypothesis. Since the initial prior is vacuous and test 1 is the strongest indicator, the chosen first test in each realization is test 1.

If proposition A is true, 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. Test 1 passes again, further confirming proposition A .

If either proposition B or C are true, test 1 fails, resulting in significant ambiguity after the first test. 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 definitive 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 .

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.

TABLE III

EXAMPLE CASE 2: RESULTS BASED ON TRUE HYPOTHESIS REALIZATION

Truth	Tests	$\text{Pr}_{\text{pl}}(A)$	$\text{Pr}_{\text{pl}}(A)$	$\text{Pr}_{\text{pl}}(C)$	$\tilde{H}_{JS}$
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

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.

C. Example Summary

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

While both cases illustrated use a single hypothesis, multiple hypotheses simply require more agent-pairs for the new propositions. The only additional requirement is user-defined priority weightings for the hypotheses, utilized in schedule incongruity resolution. The simulated results in Section V illustrate multi-hypothesis resolution through a sensor network tasking scenario.

V. SIMULATION RESULTS

This section contains a more nuanced application of JER scheduling sensor network actions to resolve multiple space situational awareness (SSA) hypotheses. SSA is concerned with accurately representing the state knowledge of objects in the space environment to provide better prediction capabilities for threats such as potential conjunction events. There are currently over 20,000 trackable objects in the space object catalog [15]. While Earth orbit is a vast volume, useful or strategic orbit regimes (e.g. low Earth orbit (LEO), Geostationary Earth Orbit (GEO)) have quickly become congested and contested [16].

Tracking techniques used in the space surveillance system still largely rely upon human-intensive models and applications [9]. As the space object population increases, the amount of data required to maintain SSA also increases [16], motivating development of autonomous sensor tasking capabilities. 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 [17], [18].

A. 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 minutes prior to the notification during a critical orbit-raising maneuver. The objective is to re-acquire the space object and diagnose the anomaly to regain situation awareness.

Anomalous GTO objects are particularly difficult to characterize as the range prohibits use of radar, requiring a wide state-space search using electro-optical sensors. Timely re-acquisition is critical since a wayward spacecraft becomes a collision risk in a densely populated orbit regimes of LEO or GEO. 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.

B. Dynamics

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 simulated using a radiometric model, including simulated effects for background sky irradiance and atmospheric transmittance [19] with cannonball-model illumination estimation. The sensor-tasking time span is limited by observation constraints (e.g. short horizon-to-horizon times in LEO, eclipse, adverse weather), placing a 15-minute time limit on the hypothesis resolution. The sensors may change actions each minute, and a receding time-horizon of two minutes is used.




Fig. 2. Test Case: propulsion explosion, schedule and entropy

C. Belief Function Models

A limited subset of potential failure modes is analyzed for illustrative purposes in this test case. The anomaly is characterized at the subsystem level to determine root-cause. The hypotheses considered include: propulsion status, navigation status, and collision in LEO. Since multiple point-of-failure events are exceedingly rare, an assumption is made that the anomaly results from a single point-of-failure.

Each hypothesis has a unique frame of discernment, used to construct JER agent-pairs. The propulsion status hypothesis consists of three propositions: nominal, non-start, and explosion. The navigation status hypothesis consists only of two propositions: nominal and anomalous. The collision in LEO hypothesis consists of $R + 1$ propositions: none, and R collision propositions (one for each potential collision object). For this study, three resident space objects (RSOs) are considered for collisions: $R = 3$. There are a total of eight JER agent-pairs in this simulation: three for propulsion failure, one for navigation status, and four for collision in LEO. Hypotheses are considered resolved if the normalized J-S entropy drops below the threshold value of $\hat{H}_{JS,thr} = 0.05$.

D. Evidence to Belief Function Mappings

Each candidate action is evaluated for its estimated effect on the hypotheses to develop evidence-to-belief-function mappings. This process is highly problem-specific, requiring the modeler to consider what each potential successful or missed detection contributes to each hypothesis. For instance, a missed detection may indicate anomaly, but if the estimated electro-optical probability of detection [19] predicted a low chance of success, that evidence is vacuous.

E. Test Case: Propulsion Explosion

In this test case, a catastrophic propulsion anomaly results in an explosion, scattering debris in the LEO parking orbit. The resulting sensor tasking schedule and hypothesis entropies are shown in Fig. 2, and resultant hypothesis resolutions (belief and plausibility) are plotted in Fig. 3




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

The first actions, tasking against GTO and LEO, result in missed detections of the primary spacecraft, contributing weak evidence toward anomalous propositions for both the propulsion and navigation statuses. However, the LEO action also successfully detects the “Coll 0” and “Coll 2” objects, refuting their collision propositions.

The sensor network initiates a search in GTO to confirm the navigation status, searching for the primary spacecraft off-nominal in GTO. During 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 (at 02:11), arriving at the correct anomaly resolution of explosion. This test case serves as a prime example of the unbiased resolution focus of JER, rejecting the incorrect prior using further evidence.

A brute-force entropy-greedy scheduler was also implemented for comparison to the complexity and bias-related improvements of 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.

An immediate difference is the number of sequence evaluations required at each iteration: 1,024 for brute-force compared to a maximum of 271 for JER (including all agent-pairs and the combination sequences). Furthermore, the incorrect potential-collision prior (induced at 02:06) is never rejected in the entropy-greedy approach: “Coll 1” is not tasked against as other actions are predicted to reduce entropy more. The entropy-greedy approach does not fully resolve the hypotheses, further underscoring the impact of the alternating-agent scheme in rejecting confirmation bias.

VI. CONCLUSION

This work develops the JER approach, a generalized evidence-gathering framework for hypothesis resolution. Employing a hypothesis abstraction enables predictive tasking and supports human cognition at a strategic and planning

level. The use of Dempster-Shafer theory to model hypothesis knowledge quantifies ambiguity and conflict. Applying a principle of equal-effort through alternating agents provides impartial hypothesis resolution to combat confirmation bias while also prioritizing time-efficient hypothesis resolution. The inclusion of efficient minimax algorithms and a hypothesis resolution pruning condition further improve computational tractability. The simulated results for a GTO insertion maneuver anomaly scenario show JER confirming or refuting propositions appropriately using evidence, avoiding fixation on incorrect propositions induced by confirmation bias.

REFERENCES

- [1] M. R. Endsley, “Situation awareness global assessment technique (SAGAT),” in *IEEE 1988 National Aerospace and Electronics Conference*, May 1988.
- [2] G. A. McIntyre and K. J. Hintz, “An information theoretic approach to sensor scheduling,” in *Proceedings of SPIE Signal Processing, Sensor Fusion, and Target Recognition V*, vol. 2755, 1996, pp. 304–312.
- [3] A. D. Jaunzemis, D. Minotra, M. J. Holzinger, K. M. Feigh, M. W. Chan, and P. P. Shenoy, “Judicial evidential reasoning for decision support applied to orbit insertion failure,” in *1st IAA Conference on Space Situational Awareness*, Orlando, FL, Nov. 2017.
- [4] B. R. Cobb and P. P. Shenoy, “On the plausibility transformation method for translating belief function models to probability models,” *International Journal of Approximate Reasoning*, vol. 41, no. 3, pp. 314–330, Apr. 2006.
- [5] D. Ellsberg, “Risk, ambiguity, and the Savage axioms,” *The Quarterly Journal of Economics*, vol. 75, no. 4, pp. 643–669, Nov. 1961.
- [6] R. Jirousek and P. P. Shenoy, “A new definition of entropy of belief functions in the Dempster-Shafer theory,” *International Journal of Approximate Reasoning*, vol. 92, no. 1, pp. 49–65, 2018.
- [7] G. Shafer, “The combination of evidence,” *International Journal on Intelligent Systems*, vol. 1, pp. 155–179, 1986.
- [8] R. R. Yager, “Arithmetic and other operations on Dempster-Shafer structures,” *International Journal of Man-Machine Studies*, vol. 25, no. 4, pp. 357–366, 1986.
- [9] F. R. Hoots and P. W. Schumacher, “History of analytical orbit modeling in the U.S. space surveillance system,” *Journal of Guidance, Control, and Dynamics*, vol. 27, pp. 174–185, 2004.
- [10] P. P. Shenoy and G. Shafer, “Axioms for probability and belief-function propagation,” in *Uncertainty in Artificial Intelligence*, R. D. Shachter, T. Levitt, J. F. Lemmer, and L. N. Kanal, Eds., vol. 4. Amsterdam: Elsevier, 1990, pp. 169–198.
- [11] D. Frey and M. Rosch, “Information seeking after decisions,” *Personality and Social Psychology Bulletin*, vol. 10, no. 1, pp. 91–98, mar 1984.
- [12] M. B. Cook and H. S. Smallman, “Human factors of the confirmation bias in intelligence analysis: Decision support from graphical evidence landscapes,” *Human Factors: The Journal of the Human Factors and Ergonomics Society*, vol. 50, no. 5, pp. 745–754, oct 2008.
- [13] E. Jonas, S. Schulz-Hardt, D. Frey, and N. Thelen, “Confirmation bias in sequential information search after preliminary decisions: An expansion of dissonance theoretical research on selective exposure to information,” *Journal of Personality and Social Psychology*, vol. 80, no. 4, pp. 557–571, Apr. 2001.
- [14] D. E. Knuth and R. W. Moore, “An analysis of alpha-beta pruning,” *Artificial Intelligence*, vol. 6, no. 4, pp. 293–326, 1975.
- [15] “Orbital debris quarterly news,” NASA Orbital Debris Program Office, Tech. Rep., Feb. 2017.
- [16] T. Blake, M. Sanchez, J. Krassner, M. Georgen, and S. Sundbeck, “Space domain awareness,” DARPA, Tech. Rep., Sep. 2012.
- [17] P. A. Brown, “Promoting the safe and responsible use of space: Toward a 21st century transparency framework,” *High Frontier*, vol. 5, no. 1, 2008.
- [18] J. D. Ianni, D. L. Aleva, and S. A. Ellis, “Overview of human-centric space situational awareness science and technology,” Air Force Research Lab, Human Performance Wing (711th), Human Effectiveness Directorate, Wright-Patterson Air Force Base, OH, Tech. Rep., 2012.
- [19] R. D. Coder and M. J. Holzinger, “Multi-objective design of optical systems for space situational awareness,” *Acta Astronautica*, 2016.