SSA Decision Support System Development and Evaluation using Cognitive Systems Engineering

Andris D. Jaunzemis, Dev Minotra, Karen M. Feigh, Marcus J. Holzinger
Georgia Institute of Technology
Moses W. Chan
Lockheed Martin Space

Abstract

Existing approaches for sensor network tasking in space situational awareness (SSA) rely on techniques from the 1950s and limited application areas while also requiring significant human-in-the-loop involvement. Increasing numbers of space objects, sensors, and decision-making needs create a demand for improved methods of gathering and fusing disparate information to resolve hypotheses about the space object environment. This work focuses on the cognitive work in SSA sensor tasking approaches. A prototype decision support system is developed using a subset of design and information relationship requirements developed through the application of cognitive systems engineering to SSA. This prototype is used in a human-in-the-loop experiment designed to compare the hypothesis resolution and situation awareness effects of hypothesis-based approaches against traditional sensor tasking approaches.

1 Introduction & Motivation

Endsley defines situation awareness as "the perception of the elements in the environment within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future," [1] or, more simply, as "knowing what is going on around you" [2]. Space situational awareness (SSA) is concerned with accurately representing the orbital state (position and velocity) of space objects "to predict, avoid, deter, operate through, recover from, and/or attribute cause to the loss and/or degradation of space capabilities and services" [3]. Currently, the space object catalog, a central database of orbit states for space objects large enough to track, contains over $20,000$ objects [4, 5], ranging from decommissioned rocket bodies to active telecommunications satellites to university science and technology experiments. The number of cataloged objects is expected to grow considerably in the near-future due to increased launch capabilities [6], improved sensor sensitivities, and continued debris generation [7]. Maintaining SSA is essential to the command and control missions of the Joint Space Operations Center (JSpOC) [8], especially as the space object environment continues to become more congested and contested [9, 10].

The space object catalog is generated and updated through sensor (e.g. telescopes, radar) observations from networks such as the Space Surveillance Network (SSN); therefore, one active research area in SSA is the allocation of sensor resources to gather evidence required to answer SSA questions in a timely manner. Current tasking techniques used in the Space Surveillance Network (SSN) rely largely on models and applications from the 1950s [11], so research increasingly focuses on improved methods for gathering and processing actionable knowledge to acquire and maintain SSA. Many of the proposed sensor tasking approaches focus on minimizing the uncertainties in the estimations of position and velocity for objects in the catalog. This process, known as space object catalog maintenance, allows improved predictions of space object motion into the future. While uncertainty in position and velocity are appropriate measures for certain SSA problems, such as collision avoidance, other decision-making questions require access to different metrics; for instance, the brightness of a satellite, measured over time, can be used to estimate the orientation and active mode of a space object [12]. Similarly, the questions decision-makers would like to answer have become more sophisticated. Previously, questions about an object focused primarily on position and trajectory, and hypotheses could easily be generated to correspond. Now, decision-makers want to know intent, operational mode, and other characteristics not necessarily confined to position or trajectory knowledge. The expansion of usable metrics and questions about satellites creates a need for more robust techniques that can represent diverse hypothesis knowledge and fuse information from a wide variety of sources to improve the overall understanding of the space object environment and answer specific SSA decision-maker questions (e.g. "is satellite X currently in a ground-observation mode?").

Real operational scenarios require a proposed sensor tasking framework and associated decision-support system to reliably support an operator's cognitive work. This paper presents the design and evaluation of a prototype decision

support system (DSS) for SSA. The prototype DSS applies design requirements derived from a cognitive work analysis (CWA) for SSA [13] and is evaluated using two different methods for scheduling sensor observations to assess the situation awareness and performance effects of both approaches.

2 Background

The following sections present background material to familiarize the reader with the SSA and cognitive engineering topics that will be discussed in this paper.

2.1 Decision Support Systems

Systems engineers face a difficult design problem in effectively managing complex sociotechnical systems, comprised of both the technological systems and the people or organizations that operate within them [14]. In the late 20th century, Keen presented a road-map for decision support system (DSS) research, focusing on utilizing emerging software tools to build semi-expert artificial intelligence systems and emphasizing the value and role of experts in DSS [15]. Shim et al. updated Keen's agenda for the 21st century, highlighting increased proficiency with technology and an associated expectation of more functionality in DSS technology [16]. Shim encouraged researchers to identify areas where tools are needed to transform uncertain and incomplete data, along with qualitative insights, into useful knowledge, along with further exploiting software and technology tools.

When considering building a DSS, the design approach for the user interface must be carefully considered. The effectiveness of a decision aid depends on relationships between the representation, the domain and associated tasks, and the characteristics of the agent [17]. Designers cannot anticipate all the possible scenarios that could arise and must therefore design displays that support effective problem solving even when novel or unanticipated scenarios are encountered [17]. To this end, it is important to frame the goal of such a design task as first helping the user to focus attention on a potentially important event and then providing integrated displays that help the user to construct a deeper understanding of the context [18].

2.2 Cognitive Systems Engineering

Militello et al. [19] define cognitive systems engineering (CSE) as "an approach to the design of technology, training, and processes intended to manage cognitive complexity in sociotechnical systems." CSE aims to provide the designer with "a realistic model of how the human functions cognitively" [20]. A multitude of CSE methodologies have emerged in recent decades to inform system design by modeling human cognitive functions [14], including cognitive work analysis [21, 22], contextual design [23], hierarchical task analysis [24], and naturalistic decision-making [25]. Unfortunately, the bulk of applications of CSE methods in industry has been limited to incorporating the insights as graphical user interface elements [14].

One of the more widely-adapted CSE frameworks is cognitive work analysis (CWA), which is a framework for establishing characteristics and constraints of the work domain [21, 22]. CWA differs from other types of work analysis by focusing more on how work may be driven by constraints imposed by the domain and less on how the work is actually accomplished [22]; in other words, the emphasis is on how the work could be done, not on how work is done or should be done [26]. It is therefore a useful design tool for a new system intended to support expert work, defined as the ability to compose a process needed for a specific task as a sequence of familiar subroutines that are useful in different contexts [27]. CWA provides an approach to "characterize the constraints that define the cognitive requirements and challenges, and the knowledge, skills, and strategies that underlie both expert performance and the error-vulnerable performance of domain practitioners" [28].

The traditional cognitive work analysis, as defined by Vicenti [22], consists of five phases or dimensions with different analysis boundaries, summarized below. The first phase, work domain analysis (WDA), analyzes relationships between the purposes, priorities, functions, and resources in the domain. The second phase, control task analysis (ConTA), analyzes activities in specific situations or tasks. The third phase, strategies analysis (SA), analyzes strategies for executing an activity. The fourth phase, social organization and cooperation analysis (SOCA), analyzes the distribution of work amongst individuals and teams, as well as communication required between these entities. The final phase, worker competencies analysis (WCA), analyzes the perceptual and cognitive capabilities and limits of humans in the domain. CWA research efforts to date predominantly implement the first two phases, leading them to be the most matured analysis techniques [26].

2.3 Description of the SSA Work Domain

Inherent in Endsley's definition of situation awareness is an understanding of what is important [2], so the design of support systems for SSA must begin with establishing goals and purposes of the work domain, which consists of numerous social and technical components. 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 (i.e. collision) events. Operators in SSA operations centers, such as the Joint Space Operations Center (JSpOC), are responsible for aggregating data on a diverse space object population, ranging from active satellites to orbital debris, and conduct analyses to predict events (e.g. conjunctions) or schedule follow-on observations to maintain a catalog of space objects. The catalog typically contains information on the orbit state (i.e. position and velocity) of the object, allowing analysts to predict its motion into the future. This data is gathered from a network of sensors (primarily telescopes and radar), many of which are controlled by independent entities, which poses difficulties in gathering timely data on specific events. Sensors in the SSN are geometrically diverse and relatively sparse, especially as compared to the over $20,000$ resident space objects (RSOs) currently tracked [5]. For instance, data collections for an object in low Earth orbit (LEO) from one ground station might only occur over two 5-minute long passes in one day due to observability constraints (e.g. line of sight, cloud cover). This leaves significant time wherein anomalies (e.g. maneuvers, on-orbit break-ups) can occur without being directly observed.

In general, the sensor tasking problem addresses how to obtain, process, and utilize information about the state of the environment [29]. The sensor tasking problem is a high-dimensional, multi-objective, mixed-integer, non-linear optimization problem, and current approaches focus on tractable sub-problems (e.g. single objectives, limited target objects, limited sensors). Potential SSA sensor tasking needs include maintaining catalogs of space object state observations [30, 31], detecting maneuvers or other anomalies [32], and estimating control modes or behavior [33, 34, 3]. These objectives are generally not complementary, especially given limited sensor resources, and the different objectives require different tasking approaches; for instance, characterization (e.g. anomaly detection) prefers many observations of a small subset of the catalog, whereas catalog maintenance prefers a diverse set of observations from as many objects as possible. Discourse and activity in SSA increasingly focus on decision-making in the presence of limited resources, uncertain information, and a contested space environment [3]. Therefore, the goal of aggregating all this data is to support SSA judgment.

2.4 Candidate Sensor Network Scheduling Approaches

This section introduces the two candidate approaches for scheduling sensor network observations that will be used throughout this paper.

2.4.1 Covariance-based Sensor Tasking

The covariance-based approach for sensor network scheduling attempts to reduce orbit state estimate uncertainty as measured by orbit state covariance. Each sensor measurement has various sources of noise (e.g. pixel width for electro-optical sensors) which contribute to state uncertainty, and the orbit state covariance matrix is a common product of fusing and filtering these measurements into a state estimate. A simple covariance-based scheduling algorithm, as shown in Eqn. (1), selects the next sensor actions to minimize the weighted-sum of state covariances:

$$\mathcal{A}^* = \arg\min_{\mathcal{A}\in\mathbb{A}} \sum_{j=1}^{N} w_j \frac{\mathrm{Tr}\mathbf{P}_j^+}{a_j} \tag{1}$$

where $\mathbb{A}$ is the set of valid action sequences at the current time step, N is the number of space objects, $\mathbf{P}_j^+$ is the estimated a posteriori state covariance of the jth space object, and a_j is the semi-major axis of the jth space object. The semi-major axis scaling factor is included to account for naturally larger uncertainties at higher-altitude orbits, preventing those objects from dominating the sensor tasking.The weighting parameters represent space object prioritization so that actions will target higher-priority assets, quickly reducing their state uncertainties and maintaining that level through follow-on observations.

The covariance-based approach is effective in addressing decision-maker needs that can be directly mapped to state estimate uncertainty (e.g. collision avoidance); however, not all decision-making needs meet this condition. Support for decision-making should provide quantifiable and timely evidence of behaviors related to specific hypotheses (e.g. threats). To support this hypothesis resolution activity, existing approaches largely focus on collecting observations to identify physical states or parameters; however, many complex hypotheses require RSO behavior prediction that takes into account other RSOs, physics knowledge, and indirect information from non-standard sources. As such, judgment is an active avenue of research in SSA, using information fusion and emerging technologies to ingest varied, sparse datasets to form a coherent story of the space environment, as well as techniques for better conveying this story to decision-makers.

2.4.2 Hypothesis-based Sensor Tasking

An alternative method for sensor network scheduling, called the hypothesis-based approach, directly quantifies decision-making questions as rigorously testable hypotheses that can be interrogated using evidence gathered from sensor observations [35, 36]. Typically a hypothesis belief structure is either converted to a probability distribution (as a Bayesian approximation of the belief structure) [37] or collapsed to a probability formulation that allows a decision-maker to place bets on the hypothesis given the available evidence using familiar Bayesian constructs [38, 39]. However, recent work has shown the utility of quantifying and incorporating ambiguity [35, 36], and evidential reasoning methods such as Dempster-Shafer theory [40] are more expressive in representing ambiguity than Bayesian approaches [41]. Hypothesis entropy reduces a complex belief structure to a single measure that represents both the ambiguity (or non-specificity) and conflict inherent in the current hypothesis knowledge; consider entropy as the hypothesis resolution uncertainty, analogous to the state covariance being the orbit state uncertainty. For further information on the specifics of Dempster-Shafer theory and hypothesis entropy, refer to [40, 41].

Consider, as an illustrative example of hypothesis-based sensor tasking, a simplified version of the Judicial Evidential Reasoning algorithm [35, 42, 36]. For this hypothesis-based approach, shown in Eqn. (2), entropy becomes the primary measurement of effectiveness, as low entropy signifies that the gathered evidence points toward a consistent and specific hypothesis resolution. When using the hypothesis-based scheduler, the next sensor actions are selected to minimize the weighted-sum of hypothesis entropies:

$$\mathcal{A}^* = \underset{\mathcal{A}\in\mathbb{A}}{\arg\min} \sum_{i=1}^{H} w_i \tilde{H}_{JS}^{+}\left(m_i\right) \tag{2}$$

where $\mathbb{A}$ is again the set of valid action sequences at the current time step, H is the number of hypotheses considered, and $\tilde{H}_{JS}\left(m_i\right)^{+}$ is the estimated a posteriori normalized Jirousek-Shenoy entropy for the ith hypothesis [41, 36, 43]. Here, the weighting parameters allow direct prioritization of decision-maker hypotheses, which automates the selection of specific space object targets toward the goal of reducing hypothesis resolution uncertainty.

3 Cognitive Work Analysis Applied to SSA

This following section summarizes methods and findings from a cognitive work analysis (CWA) applied to SSA [13]. Selecting an appropriate method for scheduling sensor networks and presenting all relevant information to a decision-maker creates a big-data and data-visualization problem. Problematically, the collected data products are often affected by adverse observation conditions, uncertainties, biases, and unobservable states that may lead to further ambiguity in evidence. Poor decision-making can have significant consequences; for instance, even small objects in LEO, which move in excess of 7 kilometers-per-second, can cause catastrophic damage to expensive and important assets.

The primary challenges of the SSA work domain can be summarized using dimensions of complexity adapted from Vicente [22]:

- **Large problem space:** High number of interacting variables.
- **Dynamic:** Constantly varying states, potentially long response times between measurement opportunities.
- **High-risk:** Errors may lead to catastrophic results.
- **Social:** Multiple organizations, with competing interests in the use of space, vying for SSA data.
- **Distributed:** Geographically disparate sensor networks and organizations.
- **Uncertainty:** Sensor bias, measurement noise, and unobservability result in probabilistic and less-than full-state knowledge.
- **Disturbances:** Operators are expected to understand anomalous behavior and bring system back within nominal conditions.
- **Automation:** Operators expected to monitor and intervene quickly and decisively in off-nominal conditions.

To derive guidelines that support decision-making and address these complexities, the first two phases of CWA (work domain analysis and control task analysis) were applied to the SSA domain [13]. The following sections summarize these findings to provide context for the DSS development.

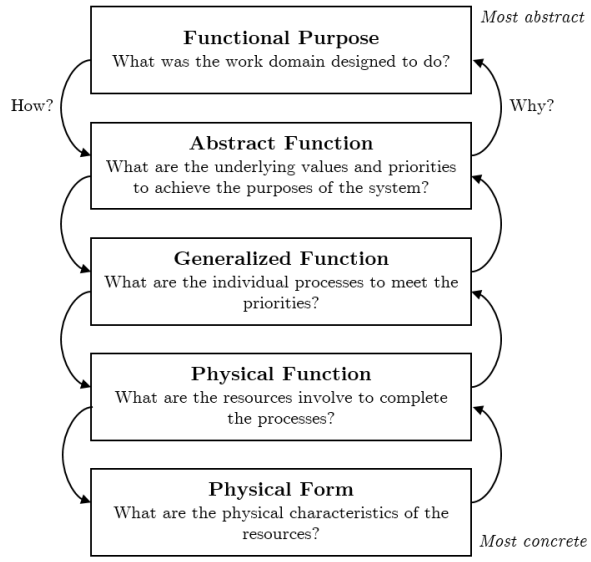

Figure 1: Abstraction hierarchy model decomposition (adapted from Naikar, 2013)

3.1 Work Domain Analysis

Work domain analysis (WDA) is employed as the first step of CWA to establish broad understandings and identify constraints that exist within the domain [14]. The result of a WDA is traditionally a model known as the abstraction hierarchy, which provides a graphical representation of linkages between purposes, priorities, functions, and resources of the domain [27, 22]. The abstraction hierarchy assesses the means-ends relationships inherent to the work domain based on the purposes of the actor(s), but explicitly does not consider the actions of the actor. In particular, the abstraction hierarchy focuses on structural means-ends relationships within the objects in the work domain. By representing the work domain from multiple levels of abstraction, an analyst can view the domain at varying levels of detail. The functional purpose and other high level abstractions provide broad overviews of the system and intended goals, whereas the lower level abstractions define attributes of the physical objects with which the actor interacts.

The structural relationships between elements in adjacent levels of the hierarchy can be summarized as follows: each level simultaneously provides a means (the "how") to elements in the level above, and an end (the "why") to the level below. The traditional five levels of decomposition included in an abstraction hierarchy [27] are summarized in Fig. 1, adapted from [44]. Beginning with the overall purposes of the work domain and progressing down the decomposition levels, the elements in the hierarchy become more concrete, arriving at the physical characteristics of the domain resources.

The following summaries provide added detail on each level of the decomposition [27, 26]:

- **Functional Purposes:** reason(s) or purpose(s) of the system.
- **Abstract Functions:** principles or priorities of the work domain that are preserved, conserved, maximized, or minimized (e.g. conservation of energy).
- **Generalized Functions:** functions that must be present for the functional purpose of the work domain to be fulfilled.
- **Physical Functions:** capabilities of the physical elements within the work domain.
- **Physical Forms:** properties of physical elements within the work domain (e.g. physical form, configuration).

It is often useful to further decompose the work domain into multiple abstraction hierarchies to separately examine different aspects of the environment and domain activities at these varying levels of detail. For instance, Burns et al. [45] used a system of three abstraction hierarchies to model naval command and control, and Miller et al. [14] used two abstraction hierarchies to separately model the environment and work domain for extra-vehicular activities (EVA)

Functional Purpose	Space Asset Safety	National Security						
Abstract Function	Orbital Dynamics	Sensor Phenomenology	Space Object Asset Priority	Hypotheses and Priorities	Workflow Efficiency			
Generalized Function	Event Detection	Conjunction Risk Assessment	*Information Fusion*	*Sensor Allocation*	Catalog Maintenance	Accuracy Degradation	Uncorrelated Track Processing	Information and Alert Dissemination
Physical Function	Computational Resources	Sensor Network Resources	Signal Transmission and Processing Capability	Personnel Capability	Complete Space Object Catalog	Individual Ephemerides	Public Space Object Catalog	
Physical Form	Number, Type, and Location of Tracked Objects	Number, Type, and Location of High-Priority Assets	Number, Type, and Location of Sensors	Signal Characteristics (uncertainty, ambiguity)				

Figure 2: Abstraction hierarchy of the SSA work domain

Abstract Function	Conservation of Mass	Conservation of Energy	Conservation of Momentum			
Generalized Function	Orbital Dynamics	Atmospheric Processes (air movement, temperature, weather)	Space Weather (radiation, solar pressure, electromagnetism)	Electromagnetic Signal Generation and Transmission		
Physical Function	Engineered Objects (signal transmission capabilities/limitations)	Natural Objects (signal transmission limitations)	Atmosphere (signal transmission capabilities/limitations)	Vacuum (signal transmission capabilities)		
Physical Form	Gravitational Field	Electromagnetic Field	Solar Radiation (luminosity, radiation)	Location and Type of Celestial Bodies (stars, planets, natural debris)	Location and Type of Man-Made Objects (spacecraft, man-made debris)	Atmospheric Conditions (turbulence, pressure, temperature, clouds)

Figure 3: Abstraction hierarchy of the SSA environment

with time-delay. Similarly, the SSA problem can be appropriately decomposed using two abstraction hierarchies: the work domain and the environment [43].

The work domain decomposition in Fig. 2, adapted from preliminary work on SSA decision-making [35], identifies capabilities and constraints that SSA operators must consider in addressing the functional purposes of maintaining space asset safety and addressing national security concerns. The environment decomposition in Fig. 3 identifies constraints external to the SSA operators, space objects, and sensor network that impact the gathering and processing of information to answer decision-maker questions. For further discussion on the elements of these decompositions, see [13].

By carefully considering the goals of SSA operators, as well as constraints imposed by the domain, this analysis develops the following insights for SSA decision-support [13]:

1. Data aggregation should account for the fusion of disparate sensor resources and various signal characteristics, including considerations of uncertainty, ambiguity, and unobservability.

2. Sensor allocation approaches should be able to directly address varied decision-maker hypotheses.

3. Fused data should be reflected through updated hypothesis knowledge states.

Applying the WDA insights to the aforementioned candidate methods for sensor network scheduling reveals that covariance-based sensor allocation and information fusion approaches, predicated on state uncertainty minimization, do not provide a robust or clear mapping to decision-maker goals of maintaining space asset safety and addressing national security concerns. Not all domain goals are readily expressed in terms of state uncertainty, and abstracting from the physical data products and state or covariance estimates to domain goals requires added cognitive effort. In contrast, applying a hypothesis-based approach provides a quantifiable means of formulating domain goals, and

resolving these hypotheses requires less abstraction from sensor tasking results to domain goals. This indicates an opportunity for improvement by allocating sensors to address specific hypotheses related to these SSA domain goals.

3.2 Control Task Analysis

A control task analysis (ConTA) is conducted to further develop insights for decision-support requirements. Previous work by the NASA directorate of Human Effectiveness has investigated how new fusion technologies could be incorporated into the SSA workflow [46, 8], including a ConTA study used to inform designs for several prototype screens for evaluation by Joint Space Operations Center (JSpOC) operators [47]. The full results of their analysis are under distribution restriction [47].

ConTA emphasizes the actions that the worker should undertake to accomplish a particular task, while also encouraging flexibility and expertise by not adhering to a strict linear, procedure-like approach. The decision ladder [27, 22], a popular choice of control task modeling, maps information processing actions and states of knowledge throughout a control task to model the cognitive processes required to complete the task.

Beginning in the bottom-left of the ladder, the *analysis* phase involves ingestion of alerts and observations to identify the system state. The *judgment* phase, at the top of the ladder, models the selection of a particular target goal through the consideration of options and their consequences. Then, descending toward the bottom-right of the ladder, the *planning* phase of the task selects actions to execute based on the stated goals, and the task terminates in the execution of that plan.

The control task decision ladders can be leveraged to generate two different types of design requirements [14, 48, 49]. A cognitive work requirement (CWR) specifies cognitive demands, tasks, and decisions that must be supported by the DSS. An information relationship requirement (IRR) specifies the context for required data, which translates that data into the actionable information that the decision-maker requires.

Miller et al. [14] demonstrate how to translate from the states of knowledge in the decision ladder to CWRs and IRRs. In their study, each state of knowledge generates at least one CWR, and each CWR has a corresponding IRR. A similar approach is followed in this work. For each applicable decision ladder state, states of knowledge articulate questions relevant to the SSA operator. Each state of knowledge generates a CWR outlining some functionality the DSS must provide to address this state of knowledge. Similarly, the CWR generates an IRR more explicitly stating the data products required to generate the necessary information.

Recalling the SSA work domain abstraction hierarchy in Fig. 2, any number of these elements may be identified for further inspection through ConTA. The full study [13] focuses on the generalized functions of information fusion and sensor allocation (highlighted in Fig. 2) to support SSA decision making. Information fusion and sensor allocation may be related in one decision ladder, as the information fusion function generally relates to the analysis phase while sensor allocation relates to judgment, planning, and execution.

In total, the ConTA of the information fusion and sensor allocation functions generated 14 cognitive work requirements and 14 corresponding information relationship requirements [13]. The prototype DSS developed in the following sections focuses on a subset of these requirements. For full details on the requirement derivation, refer to [13].

4 Development of a Prototype DSS for SSA

The CWA aided in deriving a number of DSS design insights and requirements that can be leveraged to improve decision support in SSA. A primary insight is that proposed methods should provide robust and clear mappings to decision-maker goals, indicating an opportunity for improvement by allocating sensor resources to address specific hypotheses related to these goals. The following sections focus on a subset of the design requirements derived through the ConTA, utilized in the development and validation of a prototype DSS for SSA. In particular, the design requirements selected for further analysis are ones that allow further investigation of how hypothesis-resolution supports decision-making in SSA.

4.1 Design Requirements Addressed

As discussed previously, many existing methods of sensor tasking in SSA focus on reducing orbit state uncertainty. While this tasking goal correlates well with some SSA goals (e.g. conjunction analysis where analysts try to avoid collisions between objects), not all decision-maker questions can easily be mapped to state covariance estimates. Therefore, state uncertainty minimization methods do not provide a reliable means of tasking to resolve decision-maker hypotheses related to the overarching SSA goals of space asset safety and national security.

To make connections between covariance estimates and other SSA hypotheses, current methods require decision-makers to do significant knowledge-based reasoning in the judgment phase. Specifically, when evaluating potential

courses of action, the decision-maker must consider which objects are related to high-priority hypotheses, whether those objects are available for tasking, and how much their state uncertainty can or should be reduced. Recalling Fig. 2, this requires reasoning on several different levels of abstraction: state covariances on the physical function level, sensor allocation and information fusion on the generalized function level, and hypotheses and priorities on the abstract function level. These are three very different levels of detail of the SSA problem, and requiring the decision-maker to move between them quickly (and especially iteratively, as in the judgment phase) often leads to reduced situation awareness and increased measures of workload (e.g. effort, mental demand, frustration).

Conversely, a DSS that suggests tasking assignments based directly on resolving hypotheses allows the decision-maker to remain in the abstract function level, considering trades between priorities for the different hypotheses without having to be directly concerned with the sensor allocation or state measurements. It also frees the decision-maker to spend more mental effort on formulating hypotheses to directly support the dynamic list of SSA goals. Humans out-perform automation at abstract tasks such as prioritization, so this cognitive task stands to be a better use of human decision-maker effort.

Based on the results of the CWA, we hypothesize that SSA decision-making, situation awareness, and workload are improved when using a hypothesis-based tasking algorithm as opposed to a more traditional covariance-based scheduler. In order to investigate this claim, a prototype DSS was developed that incorporated a number of the design requirements from the earlier ConTA. The following cognitive work requirements (CWRs), a subset of the full list of requirements [13], are considered in this DSS development:

- **CWR-3:** The DSS shall translate observational data into evidence.
- **CWR-6:** The DSS shall track hypothesis resolution in comparison to prescribed thresholds.
- **CWR-9:** The DSS shall provide capability for operators to adjust hypothesis priorities.
- **CWR-10:** The DSS shall assess expected hypothesis resolution based on current prioritization.
- **CWR-13:** The DSS shall generate specific actions and requests required to reach the target hypothesis resolution.

A primary function of hypothesis-based sensor tasking, applied to SSA, is to analyze candidate tasking schedules to estimate hypothesis resolution based on the current hypothesis prioritization (CWR-10). Proposed approaches, such as Judicial Evidential Reasoning (JER) [36, 42, 35], are predicated upon developing mappings from acquired data to evidence, so any evidential reasoning application must satisfactorily address CWR-3. Through the use of hypothesis entropy, quantifying both conflict and ambiguity, evidential reasoning provides a means for tracking hypothesis resolution against prescribed thresholds (CWR-6) Additionally, adjusting hypothesis weights allows decision-makers to prioritize actions that appropriately address hypotheses (CWR-9). Finally, the end result of a hypothesis-based tasking approach is the list of actions that gather the required evidence to resolve the hypotheses (CWR-13). Therefore, an evidential reasoning approach to sensor tasking, such as JER, appropriately addresses these five CWRs.

This study aims to compare a hypothesis-based tasking algorithm with a more traditional covariance-based scheduler for the purposes of supporting decision-making performance and situation awareness in SSA. The following sections describe the specifics of the prototype DSS design, including the two sensor-tasking schedulers and an overview of the functionality of the simulation environment.

4.2 Sensor Tasking Schedulers

As presented in the background, two common sensor scheduling approaches are covariance-based and hypothesis-based techniques. In the covariance-based mode, implemented as shown in Eqn. 1, the operator is able to assign specific space objects as potential actions to attempt to increase or reduce tasking actions taken on a particular object. When a given space object is prioritized, the weighting parameter for all the prioritized objects is equalized; therefore, if $n < N$ objects are prioritized, each one has a weighting of $w_i = \frac{1}{n}$ while the non-prioritized objects have a weighting of $w_i = 0$. In the hypothesis-based mode, implemented as shown in Eqn. 2, the operator is able to prioritize hypotheses directly to attempt to increase or reduce tasking actions taken against a particular hypothesis. Similar to the covariance-based mode, when a given hypothesis is prioritized, the weighting parameter for all the prioritized hypotheses is equalized. Each hypothesis may have more than one relevant space object, so prioritizing a related hypothesis does not guarantee any observations of a given object.

In both modes, the operator's assigned priorities may affect the scheduled sensor actions; however, the operator is not able to directly assign any space object or hypothesis to any sensor. Therefore, if the object's covariance is already sufficiently small or the evidence available is weak, the operator may not be able to override the algorithm.

Figure 4: Screenshot of the prototype SSA DSS, using the hypothesis-based scheduler

4.3 Simulation Environment

The final prototype SSA DSS is shown in Figs. 4 and 5. The majority of the interface is devoted to displaying relevant space object, hypothesis, and sensor network information, while the 3D view is still available but relegated to the top-right corner of the screen. The top-left panel contains information on the current simulation time and the time remaining to until the next actions are taken. The top-middle panel contains the current and proposed sensor tasking schedules, based on the operator's selections in the tasking lists.

The tasking lists on the left of both interfaces allow for drag-and-drop prioritization, while also showing current and predicted information on that object or hypothesis. All items in the "tasked" panel receive an equal weighting, while all items in the "non-tasked" panel receive a weight of zero. In the hypothesis-based mode (as shown in Fig. 4), the tasking list items display general hypothesis information and current resolution (i.e. current proposition probabilities and hypothesis entropy) and a time-history plot of entropy since the start of the simulation, along with predictions for the entropy given the next proposed tasking. In the covariance-based mode (as shown in Fig. 5), the tasking list items display space object information instead, including the current and predicted state uncertainty along with a time-history plot of state uncertainty. The drag-and-drop motion works the same for both scheduler types, only affecting which type of items receive weightings for that particular scheduler. The middle panel stores the items not relevant to that particular scheduler's tasking priorities, only displaying current information for the space objects (in hypothesis-based mode) or hypotheses (in covariance-based mode).

The status panel on the right side of the interface displays quick at-a-glance information on the sensor network, including graphical summaries of the sensor observation conditions, total space object uncertainty, and total hypothesis entropy. Below this panel, on the bottom-right of the interface, is the console that displays read-outs of any attempted actions and resultant gathered evidence.

5 Human-in-the-Loop Data Collection

The specific goal of the human-in-the-loop testing was to investigate the decision-making performance and situation awareness effects of using a hypothesis-based sensor-tasking scheduler as compared to a more traditional covariance-based scheduler. As described in the DSS design section, the test participant's only control over the tasking algorithm is through the modification of relative priorities on either space objects or hypotheses. In total, 11 graduate student participants conducted one training session and two data collection sessions each. The participants operated

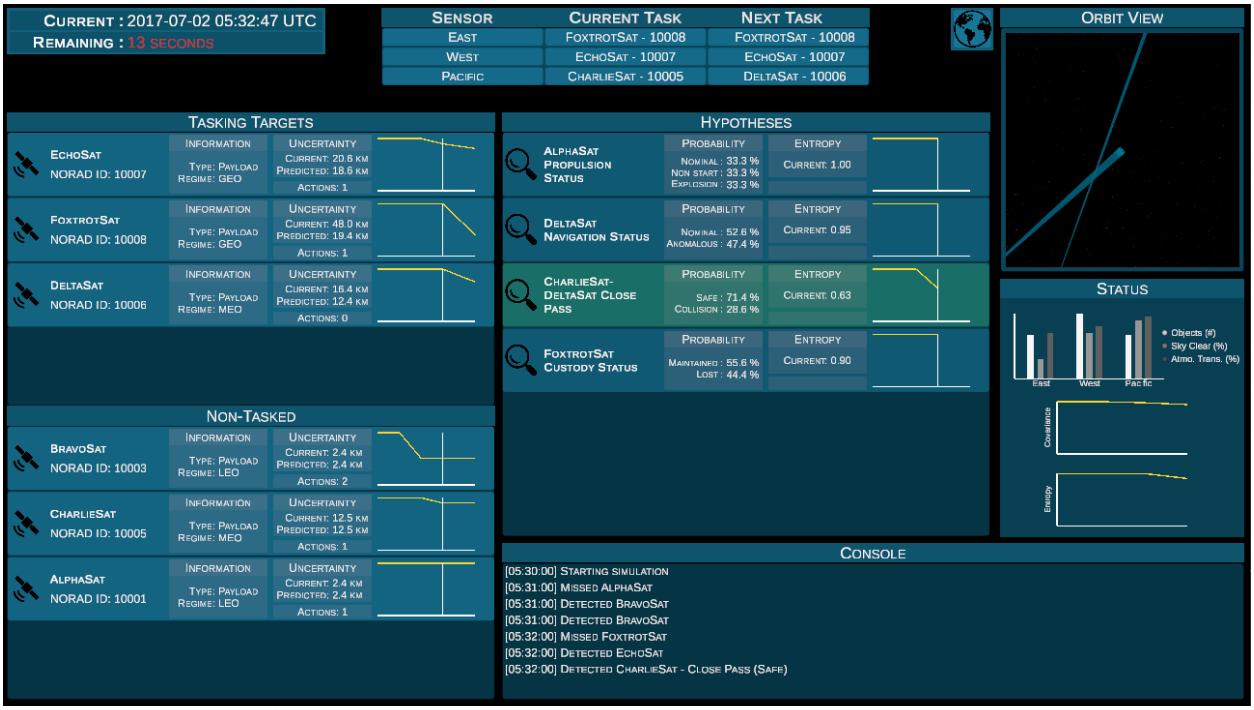

Figure 5: Screenshot of the prototype SSA DSS, using the covariance-based scheduler

Table 1: Summary of hypotheses and propositions for simulations

Hypothesis	**Propositions**
Close Pass	Safe, Collision
Propulsion Status	Nominal, Non-Start, Explosion
Navigation Status	Nominal, Anomalous
Custody Status	Maintained, Lost

each scheduler through five test scenarios for a total of 55 pair-wise data points for each dependent variable.

Each test scenario allowed a total of seven sets of actions to be scheduled and executed on the sensor network at one-minute intervals over the seven-and-a-half minute simulation. Each scenario contains three ground-based sensors with variable observation conditions, where poor observation conditions will cause observations to miss and gather zero evidence. Each scenario contains six-to-seven space objects ranging in orbit regimes from Low Earth Orbit (LEO), Medium Earth Orbit (MEO), and Geostationary Earth Orbit (GEO). Each scenario also consists of five-to-six relevant hypotheses, where the types of hypotheses and their respective propositions are summarized in Table 1.

5.1 Independent Variables

The only independent variable for this study was the scheduler type: covariance-based or hypothesis-based. Both scheduler types were implemented in the DSS described in the previous section, so the user could always see current information such as space object state uncertainty, hypothesis entropy, sensor conditions. When operating in the covariance-based mode, the scheduler also displayed predicted changes in the space object state uncertainty from the planned actions. Conversely, when operating in the hypothesis-based mode, the scheduler displayed predicted changes in hypothesis entropy from the planned actions. This allowed the user to re-prioritize space objects (in covariance-based mode) or hypotheses (in hypothesis based-mode) to achieve a desired result.

5.2 Dependent Variables

Each participant was instructed and reminded of the primary task being to answer questions (i.e. hypothesis resolution), and a number of measures were used to study the user's ability to do so, along with other relevant situation awareness factors. The data collection methods are similar to a recent DSS evaluation study published to assess emergency department information displays [50], consisting primarily of questionnaires.

5.2.1 Performance

The quality of the user's ability to answer questions is best measured through hypothesis entropy, which captures both non-specificity and conflict in the answer to the relevant question. The ideal sensor scheduling would lead to full specificity and zero conflict in the resultant hypothesis knowledge, represented by zero hypothesis entropy. Therefore, the hypothesis entropy after each set of actions was stored to allow for analysis of the evolution the hypothesis knowledge throughout the scenario. The average entropy at the end of the simulation provides the primary measure that can be used to compare hypothesis resolution performance between the scheduler types.

5.2.2 Situation Awareness

Before any questionnaire is presented, the simulation display is replaced entirely by the questionnaire so that the participant cannot rely on the DSS to answer the questions [2]. The situation awareness questionnaires were executed according to Endsley's SAGAT methodology [2], pausing for an initial questionnaire without warning during the middle of the scenario. Participants are asked ten questions relating to the scenario and hypotheses before returning to the simulation.

The specific prompts range from perception-related (Level 1) to comprehension-based (Level 2) questions. For instance, the participant might be asked which space object is related to a Custody Status hypothesis, or which sensor has the worst observation conditions, both of which can be related to perception of specific elements on the display. Alternately, the participant might be asked to compare the hypothesis resolution of a set of hypotheses and report the most- or least-resolved, or which hypothesis had the stronger evidence, both of which are more related to comprehension of the situation based on incoming information. After completion of the first stage of SAGAT questions, the simulation resumed until the end of the scenario, where ten different questions are asked. Participants were introduced to each of the potential situation awareness questions during the initial training session.

6 Results and Discussion

The data obtained from all 11 participants was compiled, and both the objective and subjective measures were statistically analyzed. Primarily, this data is analyzed through a two-factor analysis of variances (ANOVA), where the two independent variables (factors) are the scheduler type (covariance- or hypothesis-based) and the scenarios (A through E). Since the primary goal is to compare the scheduler types in terms of effectiveness, a two-factor analysis was used to determine if there was any interaction between the scheduler type and the test scenario. Using a statistical power ($1 - \beta$) of 80% and type 1 error rate (α) of 5%, the null hypothesis is considered rejected if the p-value is less than the type 1 error rate. In the results below, any statistically significant responses are marked with an asterisk, and the scheduler that provided the better performance or responses is also marked with an asterisk.

Table 2: SSA HITL results: two-way ANOVA statistics

Measure	**Covariance Mean**	**Hypothesis Mean**	$F_{1,4}$	p
*Avg. Entropy	0.24	0.23	5.97	**0.016**
SAGAT Score	15.96	15.75	0.42	0.520
SAGAT Conf.	83.46	82.33	0.47	0.494

Table 2 contains the results from this statistical analysis, with the significant results indicated by an asterisk. Though not included in the table, the interaction results for each measure were not statistically significant (all with p-values well above 0.05), indicating that the trends observed between in comparison between the scheduler types were consistent across all five scenarios. The following sections will discuss the individual components of this analysis in more detail. The full study [13] includes analysis of additional factors, including workload as measured through the NASA Task Load Index (NASA-TLX) and cognitive support questionnaire responses.

6.1 Hypothesis Resolution

Since the operator's primary goal is to resolve hypotheses, the average entropy value of all hypotheses at the end of each scenario provides an objective measurement of the participant's performance. Figure 6 summarizes the average entropy results using both the covariance- and hypothesis-based schedulers. While trends are hard to establish from this particular figure, the average entropy ANOVA result in Table 2 shows significance ($F_{1,4} = 5.97$, $p = 0.016$), with the sample mean values showing improvement using the hypothesis-based scheduler. The interaction between the scheduler types and scenarios is insignificant ($F_{1,4} = 0.39$, $p = 0.815$), indicating that the users consistently achieved lower average entropy values using the hypothesis-based scheduler regardless of the particular test scenario.

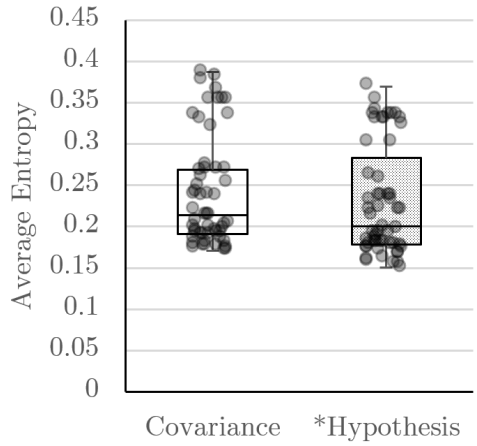

Figure 6: Average entropy vs. scheduler type

6.2 Situation Awareness

Figure 7 summarizes the responses to the SAGAT questionnaires to measure situation awareness. The score component is computed as the number of correct responses, with a maximum of 20 possible for each test case (10 mid-simulation and 10 at the end). Each situation awareness question was accompanied by a slider for the participant to indicate his/her confidence in correctly answering the question. The distributions of SAGAT responses are similar, an observation supported by the ANOVA results for both score ($F_{1,4} = 0.42$, $p = 0.520$) and confidence ($F_{1,4} = 0.47$, $p = 0.494$). Therefore, strong claims cannot be made about the impact of either scheduler on performance in answering the situation awareness questions in this study. The interaction between the scheduler types and scenarios is also insignificant ($F_{1,4} = 0.47$, $p = 0.494$). An additional ANOVA analysis was performed by separating the stage 1 (mid-simulation) and stage 2 (post-simulation) SAGAT scores to see if the trend changed throughout the simulation, but the results were nearly identical to the full SAGAT analysis shown here.

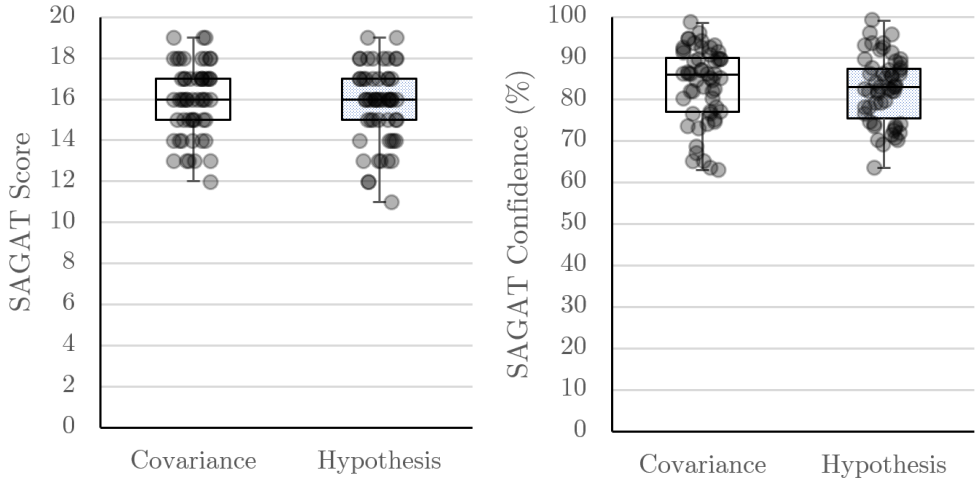

(a) Score

(b) Confidence rating

Figure 7: SAGAT responses vs. scheduler type

6.3 Limitations

While many steps were taken to ensure these tests are representative of a realistic SSA decision-making operation, the results should be interpreted carefully with appropriate caveats in light of the limited fidelity of laboratory simulation to avoid misleading or invalid conclusions. Each participant was required to have a basic level of familiarity

with orbital mechanics and spacecraft problems, but operators in SSA decision-making typically receive at least 3-6 months of in-depth training on the SSA field. Therefore, the scenarios presented to the test participants were simplified versions of realistic SSA decision-making problems, not intending to represent the full breadth of activities involved in SSA decision-making. Due to limited available training and testing time, simplified operational scenarios, and laboratory setting, these results are primarily intended to identify trends for the development of more in-depth human-in-the-loop tests closer to operational scenarios.

Perhaps the most significant caveat is the lack of a baseline comparison-point for the developed SSA DSS system. While some existing spacecraft systems engineering tools (such as the Systems Tool Kit) can be adapted to help address SSA concerns, existing software for SSA decision-making is custom-built, classified, and limited in applicability to specific use cases. The field has only recently begun focusing on information fusion and hypothesis-resolution as an important aspect of maintaining SSA, so any existing tools are likely in their infancy. Therefore, this display was developed as a prototype applicable to the simplified scenarios developed in this study and refined through pilot testing. These results may be considered a starting point for further human-in-the-loop considerations in the development of more mature SSA hypothesis-resolution DSSs, and more specific tests should be performed as these systems are developed and refined to ensure decision-maker support.

7 Conclusions

This work presents an analysis of SSA decision support from a cognitive work perspective. A subset of design requirements, developed through a cognitive work analysis applied to SSA, is used to develop a DSS driven by both covariance-based and hypothesis-based sensor tasking algorithms. The DSS developed addresses the selected design requirements by automating the translations from physical data products to evidence (CWR-3) and applying that evidence to update hypothesis knowledge (CWR-6). Additionally, while the covariance-based still displays the fused hypothesis information, the hypothesis-based approach allows the operator to directly prioritize hypotheses (CWR-9) and estimates the expected hypothesis resolution (CWR-10) from the generated set of sensor network actions (CWR-13). The DSS prototype is used in a human-in-the-loop study to assess operator performance and situation awareness. While the situation awareness results did not indicate significant differences between the scheduler types, analysis of the objective performance metric (average entropy) indicates significant improvement using the hypothesis-based approach. The full study [13] contains further analysis using metrics for workload and cognitive support. Since DSSs for hypothesis-resolution in SSA are still in their infancy, these results should be used primarily to identify trends for the development of more in-depth human-in-the-loop tests closer to operational scenarios.

References

[1] Mica R Endsley. Situation awareness global assessment technique (sagat). In *IEEE 1988 National Aerospace and Electronics Conference*, May 1988.

[2] Mica R. Endsley. *Situation Awareness Analysis and Measurement*, chapter Theoretical Underpinnings of Situation Awareness: A Critical Review, pages 3–29. Lawrence Erlbaum Associates, Publishers, Mahwah, NJ, 2000.

[3] M. Hart, Moriba Jah, David Gaylor, B. C. Ten Eyck, Eric Butcher, E. L Corral, Roberto Furfaro, E. H. Lyons, N. Merchant, M. Surdeanu, R. L. Walls, and B.J. Weiner. A new approach to space domain awareness at the university of arizona. In *NATO Symposium on Considerations for Space and Space-Enabled Capabilities in NATO Coalition Operations*, 2016.

[4] 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, June 2014.

[5] Orbital debris quarterly news. Technical Report 1, NASA Orbital Debris Program Office, February 2018.

[6] James D. Rendleman and Sarah M. Mountin. Responsible ssa coopertion to mitigate on-orbit space debris risks. In *7th International Conference on Recent Advances in Space Technologies (RAST)*, 2015.

[7] Christopher Cox, Erik J. Degraaf, Richard J. Wood, and Thomas H. Crocker. Intelligent data fusion for improved space situational awareness,. In *Space 2005*, 2005.

[8] John D. Ianni, Denise L. Aleva, and Sharon A. Ellis. Overview of human-centric space situational awareness science and technology. Technical report, Air Force Research Laboratory, Human Effectiveness Directorate, September 2012.

[9] Travis Blake, M. Sanchez, J. Krassner, M. Georgen, and S. Sundbeck. Space domain awareness. Technical report, DARPA, September 2012.

[10] T.S. Kelso. Analysis of the iridium 33 cosmos 2251 collision. In *Advanced Maui Optical and Space Surveillance Technologies*, Maui, HI, September 2009.

[11] 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*, 27:174–185, 2004.

[12] Ryan. D. Coder, Charles J. Wetterer, Kris M. Hamada, Marcus J. Holzinger, and Moriba K. Jah. Inferring active control mode of the Hubble space telescope using unresolved imagery. *Journal of Guidance, Control, and Dynamics (accepted)*, 2017.

[13] A. D. Jaunzemis, D. Minotra, K. F. Feigh, M. J. Holzinger, and M. W. Chan. Cognitive systems engineering applied to decision support in ssa. *Journal of Cognitive Engineering and Decision Making*, August 2018 (submitted).

[14] Matthew James Miller, Kerry M. McGuire, and Karen M. Feigh. Decision support system requirements definition for human extravehicular activity based on cognitive work analysis. *Journal of Cognitive Engineering and Decision Making*, 11(2):136–165, 2017.

[15] P. Keen. Decision support systems: The next decade. *Decision Support Systems*, 3(3):253–265, 1987.

[16] J. P. Shim, Merrill Warkentin, James F. Courney, Daniel J. Power, Ramesh Sharda, and Christer Carlsson. Past, present, and future of decision support technology. *Decision Support Systems*, 33:111–126, 2002.

[17] Philip J. Smith, Kevin B. Bennett, and R. Brian Stone. Representation aiding to support performance on problem-solving tasks. *Reviews of Human Factors and Ergonomics*, 2(1):74–108, 2006.

[18] Philip J. Smith, Dustin Johnson, Shawn Pruchnicki, Josh Schimmel, Amy Spencer, and Seth Young. Cognitive engineering considerations in the development of an information retrieval system: Avoid fixation on technological substitutes. *Journal of Cognitive Engineering and Decision Making*, 10(2):126–130, June 2016.

[19] Laura G. Militello, Cynthia O. Dominguez, Gavan Lintern, and Gary Klein. The role of cognitive systems engineering in the systems engineering design process. *Journal of Systems Engineering*, 13(3):261–273, 2010.

[20] E. Hollnagel and D. D. Woods. Cognitive systems engineering: New wine in new bottles. *International Journal of Man-Machine Studies*, 18:583–600, 1983.

[21] J. Rasmussen and A. M. Pejtersen. *Cognitive Systems Engineering*. Wiley, New York, 1994.

[22] Kim J. Vicente. *Cognitive Work Analysis: Toward Safe, Productive, and Healthy Computer-Based Work*. CRC Press, 1999.

[23] H. Beyer and K. Holtzblatt. *Contextual Design: Defining Customer-Centered Systems*. Morgan Kauffman Kluwer Academic Publishers, San Francisco, CA, 1998.

[24] J. Annett. *Handbook of Cognitive Task Design*, chapter Hierarchical Task Analysis, pages 17–36. CRC Press, Boca Raton, FL, 2003.

[25] Gary Klein. Naturalistic decision making. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 50(3):456–460, 2008.

[26] Matthew James Miller. *Decision Support System Development for Human Extravehicular Activity*. PhD thesis, Georgia Institute of Technology, December 2017.

[27] Jens Rasmussen. The role of hierarchical knowledge representation in decision and system management. *IEEE Transactions on Systems Man and Cybernetics*, 15(2):234–243, 1985.

[28] A. Bisantz and E. M. Roth. Analysis of cognitive framework. *Reviews of Human Factors and Ergonomics*, 3(1):1–43, November 2008.

[29] Gregory A. McIntyre and Kenneth J. Hintz. An information theoretic approach to sensor scheduling. In *Proceedings of SPIE Signal Processing, Sensor Fusion, and Target Recognition V*, volume 2755, pages 304–312, 1996.

[30] Kyle J. DeMars, Islam I. Hussein, Carolin Frueh, Moriba K. Jah, and R. Scott Erwin. Multiple-object space surveillance tracking using finite-set statistics. *Journal of Guidance, Control, and Dynamics*, pages 1–16, March 2015.

[31] Tyler A. Hobson. *Sensor Management for Enhanced Catalogue Maintenance of Resident Space Objects*. PhD thesis, University of Queensland Australia, 2014.

[32] A. D. Jaunzemis, M. V. Mathew, and M. J. Holzinger. Control cost and mahalanobis distance binary hypothesis testing for spacecraft maneuver detection. *Journal of Guidance, Control, and Dynamics*, 39(9):2058–2072, September 2016.

[33] Marcus J Holzinger, Charles J Wetterer, K. Kim Luu, Chris Sabol, and Kris Hamada. Photometric attitude estimation for agile space objects with shape uncertainty. *Journal of Guidance, Control, and Dynamics*, 37(3):921–932, May 2014.

[34] Ryan D. Coder, Marcus J. Holzinger, and Richard Linares. 3dof estimation of agile space objects using marginalized particle filters. *Journal of Guidance, Control, and Dynamics*, Submitted 2016.

[35] 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, November 2017.

[36] A. D. Jaunzemis, M. J. Holzinger, M. W. Chan, and P.P. Shenoy. Evidence gathering for hypothesis resolution using evidential reasoning. *Journal of Information Fusion (submitted, in revision)*, 2018.

[37] Barry .R. Cobb and Prakash P. Shenoy. On the plausibility transformation method for translating belief function models to probability models. *International Journal of Approximate Reasoning*, 41(3):314–330, April 2006.

[38] Bjornar Tessem. Approximations for efficient computation in the theory of evidence. *Artificial Intelligence*, 61, 1993.

[39] Philippe Smets and Robert Kennes. The transferable belief model. *Artificial Intelligence*, 66(2):191–234, 1994.

[40] Glenn Shafer. The combination of evidence. *International Journal on Intelligent Systems*, 1:155–179, 1986.

[41] Radim Jirousek and Prakash P. Shenoy. A new definition of entropy of belief functions in the Dempster-Shafer theory. *International Journal of Approximate Reasoning*, 92(1):49–65, 2018.

[42] A. D. Jaunzemis, M. J. Holzinger, M. W. Chan, and P.P. Shenoy. Evidence gathering for hypothesis resolution using judicial evidential reasoning. In *21sts International Conference on Information Fusion*, Cambridge, UK, July 2018.

[43] A. D. Jaunzemis. *Predictive Sensor Tasking and Decision Support for Space Situational Awareness using Evidential Reasoning*. PhD thesis, Georgia Institute of Technology, May 2018.

[44] Neelam Naikar. *Work Domain Analysis: Concepts, Guidelines, and Cases*. CRC Press, 2013.

[45] C. M. Burns, D. J. Bryant, and B. A. Chalmers. Boundary, purpose, and values in work-domain models: Models of naval command and control. *Systems, Man, and Cybernetics, Part A: Systems and Humans*, 35(5):603–616, 2005.

[46] D. Aleva, J. Ianni, and V. Schmidt. Space situation awareness human effectiveness research trends. In *World Congress in Computer Science, Computer Engineering and Applied Computing*, 2012.

[47] Denise Aleva and James McCracken. Jspoc cognitive task analysis. In *Advanced Maui Optical and Space Surveillance Technologies Conference*, Maui, HI, September 2009.

[48] S. S. Potter, J. W. Gualtieri, and W. C. Elm. *Handbook for Cognitive Task Design*, chapter Case Studies: Applied Cognitive Work Analysis in the Design of Innovative Decision Support, pages 1–32. Lawrence Erlbaum Associates, London, UK, 2003.

[49] W. C. Elm, S. S. Potter, J. W. Gualtieri, E. M. Roth, and J. R. Easter. *Handbook for Cognitive Task Design*, chapter Applied Cognitive Work Analysis: A Pragmatic Methodology for Designing Revolutionary Cognitive Affordances, pages 357–382. Lawrence Erlbaum Associates, London, UK, 2003.

[50] Nicolette McGeorge, Sudeep Hedge, Rebecca L. Berg, Theresa K. Guarrera-Schick, David T. LaVergne, Sabrina N. Casucci, A. Zachary Hettinger, Lindsey N. Clark, Li Lin, Rollin J. Fairbanks, Natalie C. Benda, Longsheng Sun, Robert L. Wears, Shawna Perry, and Ann Bisantz. Assessment of innovative emergency department information displays in a clinical simulation center. *Journal of Cognitive Engineering and Decision Making*, 9(4):329–346, December 2015.