Received February 20, 2017, accepted April 18, 2017, date of publication May 23, 2017, date of current version July 24, 2017.

Digital Object Identifier 10.1109/ACCESS.2017.2707091

A Context-Driven Framework for Proactive Decision Support With Applications

MANISHA MISHRA[1], DAVID SIDOTI[1], GOPI VINOD AVVARI[1], PUJITHA MANNARU[1], DIEGO FERNANDO MARTÍNEZ AYALA[1], KRISHNA R. PATTIPATI[1], (Fellow, IEEE), AND DAVID L. KLEINMAN[2], (Fellow, IEEE)

[1]Department of Electrical and Computer Engineering, University of Connecticut, Storrs, CT 06269, USA
[2]Naval Postgraduate School, Monterey, CA 93943, USA

Corresponding author: Manisha Mishra (manisha.mishra@uconn.edu)

This work was supported by the U.S. Office of Naval Research under Grant N00014-12-1-0238 and Grant N00014-16-1-2036.

ABSTRACT Major challenges anticipated in the future C^{4}ISR (command, control, communications, computers, intelligence, surveillance, and reconnaissance) operations involve rapid mission planning/re-planning in highly dynamic, asymmetric, unpredictable, and network-centric environments. Developing decision support for such complex mission environments requires automated processing, interpretation, and development of proactive decisions using large volumes of structured, unstructured, and semi-structured data, while simultaneously decreasing the time necessary to arrive at a decision. To overcome this data deluge, there is a need for mastering information dominance via acquisition, fusion, and transfer of the *right* data/information/knowledge from the *right* sources in the *right* mission context to the *right* decision-maker (DM) at the *right* time for the *right* purpose (6R). The fundamental challenge in achieving the 6R is to conceive a generic framework that encompasses the dynamics of relevant contextual elements, their interdependence and correlation to the current and evolving situation, while taking into account the cognitive status of the DM. In this paper, we propose a context-driven proactive decision support (PDS) framework that comprises: 1) adaptive model-based dynamic graph models (e.g., Dynamic Hierarchical Bayesian Networks) and the concomitant inference algorithms for context representation, inference, and forecasting, 2) information selection, valuation, and prioritization methods for context-driven operations, including uncertainty management approaches, and 3) application of PDS concepts for courses of action recommendations across representative maritime operations.

INDEX TERMS Context-aware decision support, context representation, uncertainty management, proactive decision support.

I. INTRODUCTION

A. MOTIVATION

Future C^{4}ISR (Command, Control, Communications, Computers, Intelligence, Surveillance and Reconnaissance) environments are anticipated to be more complex, distributed and network-centric due to three trends in operations faced by mission planners. First, the ubiquitous use of complex cyber-physical systems (e.g., unmanned subsurface, surface, air and ground vehicles), provides a unique range of decision options, such as ultra-long endurance and high-risk mission acceptance, which cannot be reasonably performed by manned systems. Additionally, the use of unmanned systems, combined with smart sensing technologies, provides real-time access to data, especially in regions which cannot be accessed by humans. With the availability of sensor data, making real-time decisions for efficient control, de-confliction and coordination of heterogeneous manned/unmanned resources within a congested mission space becomes immensely challenging.

Second, the changing patterns of potential threats and conflicts in today's world requires proactive military capabilities to execute a full range of operations – from normal peacetime operations (e.g., humanitarian assistance/disaster relief, search and rescue operations) to major combat operations (e.g., theater anti-submarine warfare (TASW), counter-smuggling operations). With increased networking capabilities, the new operational planning concepts emphasize network-centric distributed planning capabilities and decentralized execution of multiple simultaneous tasks, often under disrupted, intermittent and low bandwidth environments (e.g., anti-access/area denial (A2AD) for regional access and freedom of navigation) [1]. The A2AD concept is predicated on the use of stealthy submarines in littoral environments, equipped with surface-to-air missiles, and fighter planes to destroy land-based targets, while disrupting space-based sensors and communications. In this vein, coordinated, dynamic and adaptive decision support technologies are vital to counter A2AD strategies for achieving sea control, power projection, and deterrence.

Third, with the availability of real-time data, it becomes difficult to recognize and extract mission relevant information to present to the Decision Makers (DMs) in an easily understandable format, especially for time-critical operational decision making. Presenting too much information overwhelms the DM, as it requires the DM to infer or reconstruct courses of action (COAs) by comprehending all of the data, which is often laborious and an error-prone process. Recently, channelized attention (or cognitive tunneling) has been implicated in numerous operational mishaps, which could have been preventable [2]. For example, operators flying an MQ-1B Predator in Afghanistan in 2009 were so focused on a fierce firefight that they failed to notice that the unmanned aircraft was headed toward a mountain; the aircraft, valued at 3.9 million USD, was destroyed on impact.

Moreover, eight of the soldiers who were to be provided air support by the Predator were killed. If alerts and tasking had been appropriately issued to the right operators on the team, those lives may have been saved [3]. Thus, these trends suggest that there is a need to develop methods to identify, process and integrate decision-relevant information from structured, semi-structured and unstructured data and proactively present this information to DMs in order to compress the detect-to-engage timeline for anticipatory decision making.

Effective mission planning involves Boethius' *who*, *what*, *why*, *how*, *where*, *when*, *with what*, implying *who* has the expertise to make the plan (DMs who may be humans or autonomous agents), *what* needs to be planned (tasks, jobs, and actions to be executed using assets or resources), *why* make the plan (desired goal or objective function), *how* to achieve the expected outcome (the assignment of assets to tasks, sequencing of activities arranged as a directed graph), *where* the plan is executed (task location or mission area), *when* the plan is to be executed (start time and duration for each task), and *with what* facilities to make the plan (decision support systems that exploit relevant information about tasks, assets, desired objectives, DM's cognitive states, etc.) [4]. The current state-of-the-art includes the Battlespace on Demand (BonD) framework, which provides a systematic approach to convert the knowledge of forecasted environmental data into actionable decisions. The BonD framework includes four tiers: Data (Tier 0), Environment (Tier 1), Performance (Tier 2), and Decision (Tier 3) [5], [6]. Currently, the Tier 3, Decision layer, is primarily a manual capability and requires experienced and highly-trained personnel for superior mission performance. As an illustration, the TASW commanders spend many hours manually assigning submarines to waterspace for training exercises. This process requires an enormous amount of human supervision and is therefore error-prone and cumbersome. As per the US Navy's report on Fleet Battle Experiment Kilo, to address the currently slow and manually intense apportionment process, the anti-submarine warfare (ASW) commanders need training, more staff for crisis action planning at the operational level, improved collaboration among commanders, and decision support for allocating assets to engage hostile submarines and prevent attacks [7]. Therefore, there is a need to capture the evolving dynamics among the mission planning elements and provide proactive decision support (PDS) to the DMs across a range of maritime operations (e.g., counter-smuggling, ASW, waterspace planning) and across different skill levels of command personnel.

In this vein, it becomes critical to systematically identify and make appropriate inquiries regarding the key planning events as they emerge, as shown in Figure 1 (e.g., What is known? What events have happened? Why they happened? What could happen? What COAs need to be taken? What do the COAs mean?). Once these inquiries are made, selection, fusion, and transfer of the *right* data/information/knowledge from the *right* sources in the *right* context to the *right* DM at the *right* time for the *right* purpose (6R) is possible [8]. The 6R process facilitates the processing of collected data and presenting decision-relevant information to DMs in a timely manner to aid in effective decision making, even under dynamic, uncertain and challenging mission conditions (e.g., changes in mission goals, environment, assets and mission tasks/threats). By context-driven we mean dynamically integrated knowledge that is (i) relevant to the mission, the environment, assets, threats/tasks including the DM activities, (ii) informed by up-to-date data sources, and (iii) congruent with the workflow and individual DM's role in the mission, workload, time pressure and expertise. To be effective, the new mission planning concepts must be accompanied by a proactive C^4ISR decision framework that facilitates context-dependent high value data to be identified, prioritized, processed and exchanged among human DMs and automated agents based on operational information requirements. The term proactive extends beyond adaptability; its virtues range from responsiveness, robustness, innovativeness, flexibility, and anticipation of changes in the mission context.

B. LITERATURE REVIEW

Context has been extensively studied in the last decade; however, it is a poorly used source of information in decision making. In order to successfully incorporate context into the decision making process, various attempts have been made to define context and its architectural framework. In [9]–[11], context is defined, via an example, as location, environment, identity of people and objects, including the changes to them. References [12]–[14] view context as the state of the

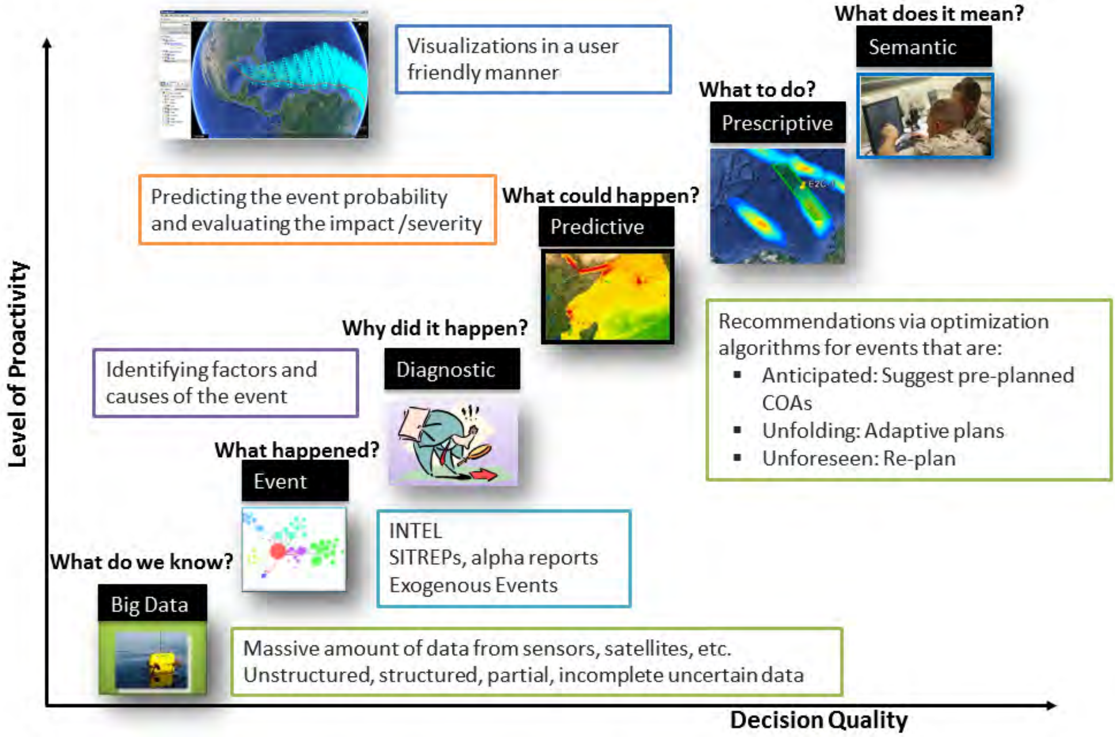

FIGURE 1. **Proactive decision support (PDS) problem involving the processing of massive sensor data to provide proactive course of action (COA) recommendations to DMs in a meaningful manner. The key is to identify the critical events and predict the event probability to evaluate the impact on the mission. Based on anticipated, unfolding or unforeseen events, appropriate COA recommendations are provided via optimization algorithms. Here, higher level of proactivity is expected to result in better decision quality.**

application's surroundings or settings. Reference [15] claims that the important aspects of context are: where the user is, who the user is with, and what resources are nearby. They define context to be a constantly changing execution environment, which includes computing, physical and user environments. Reference [16] defined context as any information that can be used to characterize the situation of entities (i.e., whether a person, place or object) that are considered relevant to the interaction between a user and an application, including the user and the application themselves. According to [16], context is typically the location, identity and state of people, groups, and computational and physical objects. In most of the literature, the authors either consider context as the physical environment (real world), computing environment, or the user environment.

Defining context by example is difficult to incorporate in a systematic decision making process because when a new context arises, it is hard to classify it correctly. Additionally, since context is about the whole situation relevant to an application and its set of users, it is difficult to enumerate all of its aspects, as these change depending on the situation [17]. Due to this diversity of contextual definitions, it is vital to categorize it in order to comprehend it in a systematic manner. Therefore, from the perspective of mission planning, understanding context at each level of the decision making paradigm of transforming data to decisions requires: i) context in the ontology of the BonD framework; ii) context in computing and decision making; and iii) context in communication, interpretation and visualization. We analyze the state-of-the-art in the above three categories, which are essential in the different tiers of the decision framework.

1) CONTEXT IN ONTOLOGY

Context facilitates in defining which knowledge should be considered, what are its conditions of activation and limits of its validity, and when to use it [18]. This is especially important for the creation and use of large and reliable knowledge-based systems for complex mission planning. Contexts act as adjustable filters for giving the right meaning in the current situation and to present the minimal number of information pieces and essential functions that are necessary to the task at hand [18]. Context representation via ontology facilitates knowledge sharing and reuse in an open and dynamic distributed system [19]–[22]. Also, it derives newly acquired knowledge and facts using reasoning on contextual data and information by using inference engines. A shared context is referred to as an ontology because the domain provides a common understanding of the involved design concepts and of the topological relations among them.

2) CONTEXT IN COMPUTING, DECISION MAKING AND COMMUNICATION

Once the context is sensed and represented via ontology, it can be instantiated for decision making. Reference [23] mentions two ways of using context by either automatically adapting the behavior according to the discovered context (active context), or presenting the context to the user in real-time and/or storing it for later retrieval (passive context). Active context aware computing is challenging, as it requires a decision support system to switch models based on the context. [18] considered context as an attribute of the interactions among agents, as opposed to context being a fixed attribute of a particular problem or application domain, i.e., without interacting agents, there would be no context. In communication, context is considered as the history of all that has occurred over a period of time, the overall state of knowledge of the participating agents at a given moment, along with the small set of things they are attending to at that particular moment. Context appears as a shared space of knowledge. Each entity involved in an interaction has its own context, which may or may not be consistent with parts of the context of other entities. Context can be thought of as a system that would be expert at 'predicting' what the user would likely want/need to do next because of its knowledge of what had happened to either that user or other users with the same goals/needs. Then, context can be provided through a user interface using currently known graphical techniques.

In the domain of decision making for mission planning, we define context as an interlinked multidimensional dynamic feature space, where a significant change in any of the elements (i.e., contextual information) incurs a corrective measure to be taken in order to obtain the desired objectives/ or stay within the acceptable performance limits of the operation. By this definition, it is crucial to note that context is proactive and not reactive; this implies that a corrective action is taken only when the changes in the mission elements are beyond the acceptable limits. Deciding when a particular change or context becomes significant enough that a corrective measure should be taken is an important research question in itself. According to [24], context-aware applications require i) heterogeneity and mobility to represent the large contextual data; ii) relationships and dependencies among the different elements of context must be captured to ensure correct behavior of the application; iii) timeliness ensures that the context aware application must have access to the past and forecasted future contextual information; iv) imperfection regarding contextual information must be taken into account (e.g., uncertainty in context); v) reasoning is necessary to determine if a particular change requires an adaptation to the new context; vi) usability would allow the system to check for "what-if" scenarios. In this paper, we focus on defining a dynamic multi-dimensional feature space and how to utilize it to develop a proactive decision support framework that is applicable across various mission planning domains.

C. SCOPE AND ORGANIZATION OF THE PAPER

The goal of proactive decision making is to identify the 6R, i.e., to determine and understand the current mission context and, acquire the concomitant decision-relevant high-value information for anticipating and exploring alternative COAs to achieve the DM's intent in a timely manner. The technical challenges in developing PDS include:

1) How to define, represent, identify and characterize critical mission context elements for dynamic resource management, so that the DM readily understands what is known and unknown about the mission scenario, and what controls/ options might be available to manage and optimize decisions given these uncertain events?
2) How to include realistic asset modeling parameters (e.g., speed, capability, sweep width) under evolving mission context, without making the problem computationally expensive?
3) How to predict context and develop flexible (for anticipated events), as well as agile (for unanticipated events) COAs to achieve resilience in dynamic and uncertain environments?
4) How to communicate context among the DMs in a timely manner?
5) How to explore methods for instantiating alternative COAs in a manner consistent with the needs of military planning/re-planning when changes in mission context are detected and the root cause inferred?
6) How many candidate COAs should the DM be presented with? If more than one, how to present them to the DM, what is the planning horizon, and what meta-information is required?
7) How frequently should re-planning occur (e.g., with a predefined frequency (few hours, daily), as events emerge, or a hybrid of the two)?
8) How to include the ability to apply weights on objectives and how to conduct "what-if" experiments based on the choice of weightings and confidence in the information?
9) How to reduce the DM's/operator's cognitive workload?

Our work in this paper goes beyond previous research on context-driven decision support by developing a broadly applicable framework for PDS and by addressing the following sub-problems. In section II, we focus on developing: i) generic and widely applicable methods for representing missions, environment, assets, threats and humans, and instantiating these models with operational data in order to detect incipient context changes early, infer the current context, project the impact of changed context on the mission goals, and proactively explore decision alternatives to exploit opportunities or mitigate the negative consequences of a changed context to achieve the DM's intent; ii) reliable methods for uncertainty management, information selection, valuation and prioritization; and iii) context dissemination and visualization in a user friendly manner. In section III, we implement the proposed context-driven PDS algorithms within different maritime domains for impact analysis and COA recommendations. We envision the context-driven decision making process to dynamically invoke plans as a function of emerging events, readily adapt plans to meet unfolding

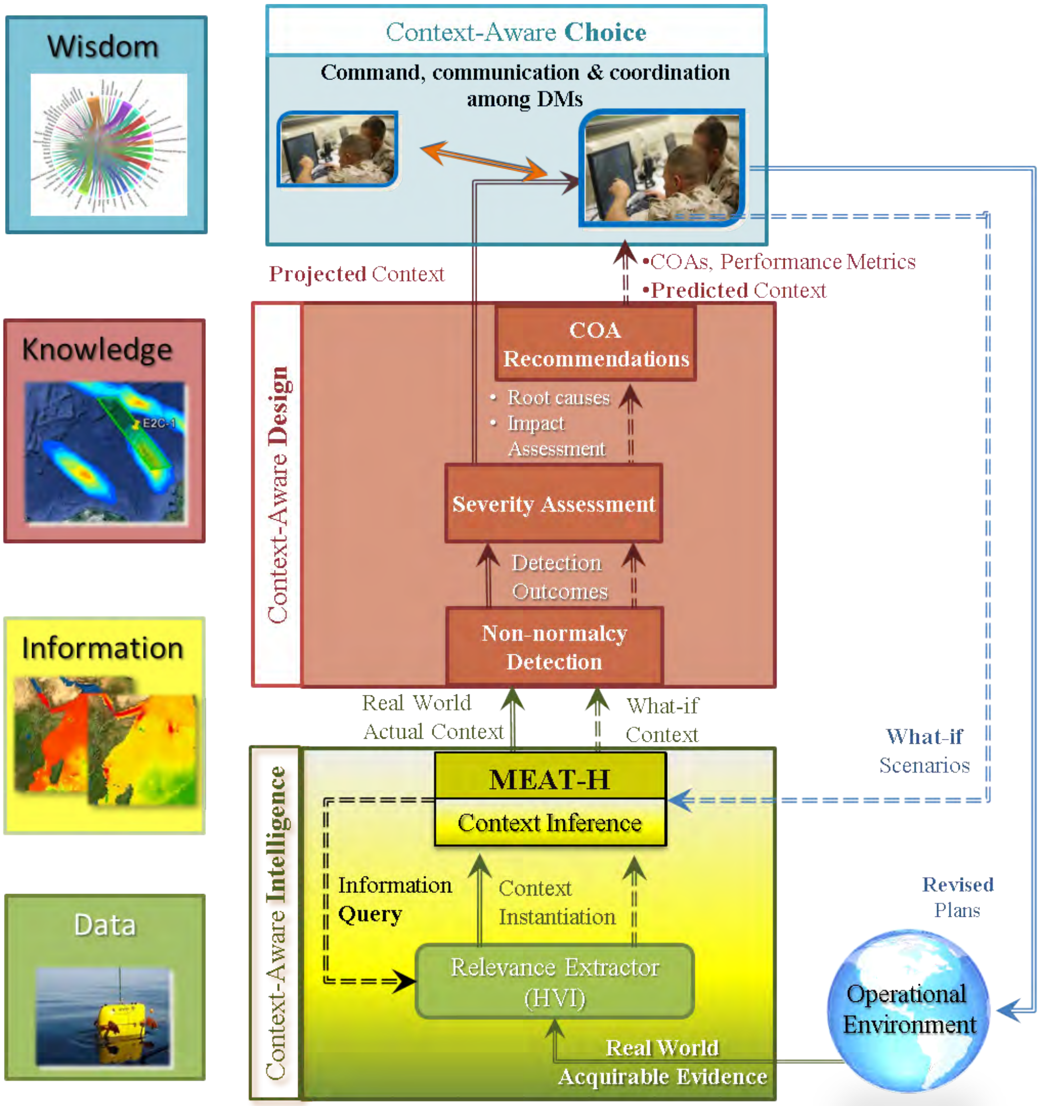

FIGURE 2. Proactive decision support (PDS) framework consists of three blocks (a) Context-Aware Intelligence where all the data collected from the environment is prioritized, filtered and arranged in the form of MEAT-H ontology to facilitate context inference; (b) Context-Aware Design allows to further comprehend the context via non-normalcy detection and severity assessment. A set of COA recommendations and the associated risk is developed by predicting the context; and (c) Context-Aware Choice allows visualization of the different COAs in an easily interpretable manner. The framework also allows the DM to tweak the parameters and conduct "what-if" analyses (denoted by dashed arrows) before finalizing the operational decisions. The solid arrows represent the flow of information/actions within the decision framework. The proposed PDS framework maps well to the existing BonD framework.

events, monitor the outcomes of many of its previous decisions, and replan if warranted, while providing transparent and unobtrusive recommendations relevant to the evolving mission scenario. Finally, section IV concludes our research work and provides future research directions.

II. PROACTIVE DECISION SUPPORT FRAMEWORK

Our PDS framework is illustrated in Figure 2 and is consistent with Simon's model of the decision making process [25], as well as Endsley's situation awareness model [26]. The PDS framework encompasses three key building blocks, as shown in Figure 2, viz., context-aware intelligence, context-aware design and context-aware choice. The context-aware intelligence block involves data collection, formatting, and filtering to develop a list of contextual elements and workflows for maritime missions (e.g., ASW, counter-smuggling). We posit context as a two-level (i.e., conceptual and operational), multi-dimensional feature space composed of Mission, Environment, Assets, Threats, and Humans (MEAT-H), that adapts to the DM's role, workload, time pressure, task expertise, and a variety of other cognitive states pertaining to human performance. Additionally, we develop the structure of the corresponding dynamic hierarchical Bayesian network (DHBN),

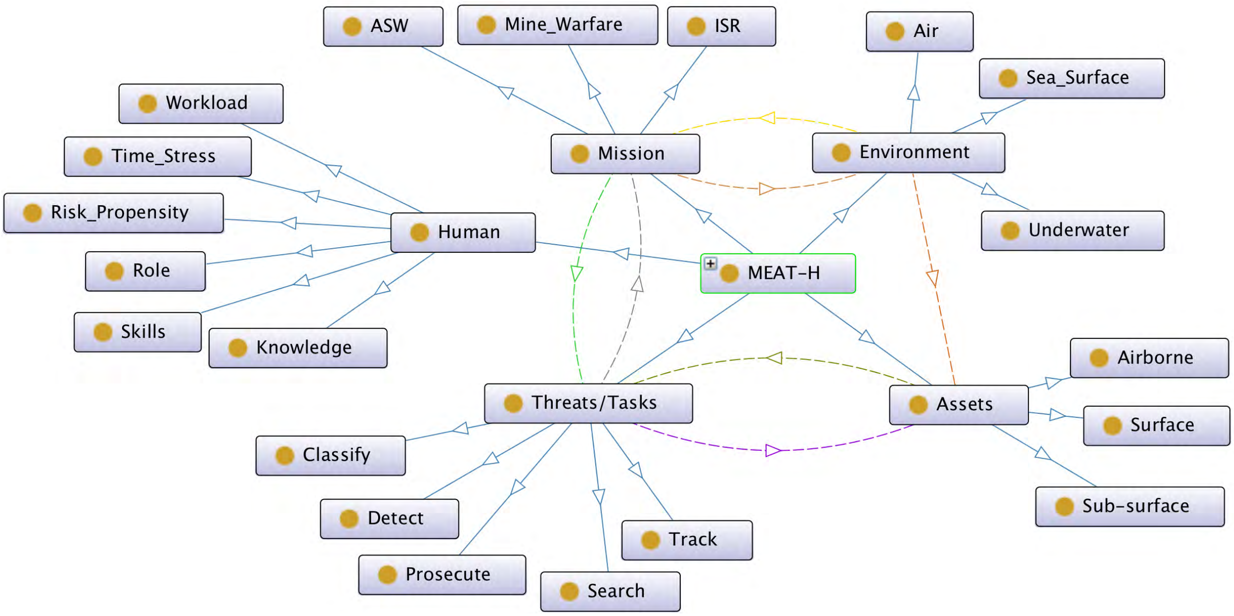

FIGURE 3. Representation of MEAT-H elements via the Protégé software [31]. The solid arrows represent the hierarchy, whereas the dashed arrows denote the relationships among different elements.

representing the abstract contextual elements of the MEAT-H architecture with the help of idioms and subject-matter experts. In the results section of the paper, we demonstrate realistic mission scenarios, where the evolving context forms a stimulus that automatically triggers the transmission of context, from the Intelligence block to the Design block in Figure 2. In the context-aware design phase, the contextual information, enriched with the non-normalcy detection and severity assessment features, triggers automatic generation of COA recommendations using optimization-based resource management algorithms across multiple scenarios. Within each of these scenarios, we simulate multiple environmental conditions such that the recommended COAs are robust even in the face of an evolving context [27]. In the context-aware Choice block, we model, capture and analyze various aspects of the DM's behavior under different contexts (e.g., varying workload, time constraints, different environmental conditions) via the internal cognitive context inference algorithms. Mission workload measures include the task load, DM's cognitive workload, as well as asset coordination [28]–[30]. The mission performance measures include task execution effectiveness, task processing delays, and coordination overheads. The proposed PDS framework represents, identifies, exploits and communicates context-relevant information to DMs with real-time context-driven COA recommendations, while evaluating "what-if" scenarios for achieving superior mission performance and mitigating error-prone planning processes. We discuss the algorithms in each of the blocks in detail in the following subsections.

A. CONTEXT REPRESENTATION

In order to successfully incorporate context within the decision-making process, it is imperative to systematically understand, identify and project operational context. The key technical challenges here are:

1) What are the methods to represent the multi-dimensional context feature space?
2) How to infer context from operational data of dynamic interactions of contextual elements?
3) How to detect non-normal situations and project the potential paths of the mission given an inferred context?

In this paper, the multi-dimensional operational context comprises an external context in terms of i) Mission, ii) Environment, iii) Assets and iv) Threats, and an internal context composed of the human/DM's cognitive states, as shown in Figure 3, via the Protégé tool [31]. Each of these MEAT-H elements may have several sub-elements with associated states and the level of specificity may vary with the role of the DM. Some examples of Mission include ASW, counter-smuggling operations, and Unmanned Aerial Vehicle (UAV) coordination in a heterogeneous environment. Each mission is characterized by goals, desired performance, achievable performance, and constraints. Examples of Environment elements include cloud cover, sea state, precipitation, salinity, and temperature, each of which may have multiple states. Examples of Asset types include sea, air, space and land assets with sub-elements such as frigates, high or medium endurance cutters,and P-3s, where the states of the sub-elements may include the asset's availability (available, unavailable) or crew endurance (measured in hours). Lastly, it is important to note that Threats in the MEAT-H framework may refer to not only the threats or the tasks to be done, but also the DM's activities to accomplish the tasks (viz., the workflow). Some examples of Threats/Tasks include interdiction of a drug smuggler, protection of a high priority maritime vessel, or reconnaissance, depending on

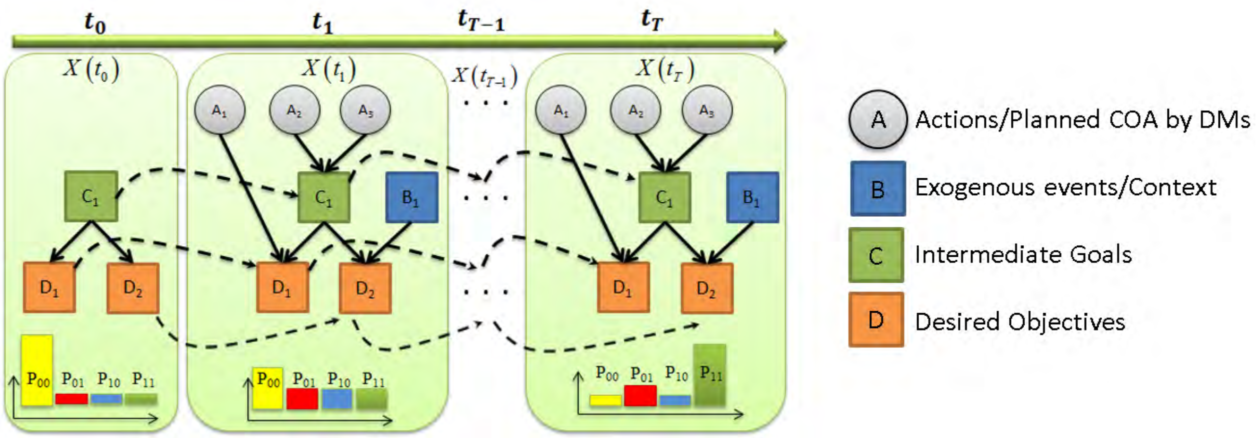

FIGURE 4. Time evolution of context-based mission using Dynamic Bayesian Networks. Actions, exogenous events/context, intermediate goals and desired objectives are represented by *A*, *B*, *C* and *D* respectively. The solid arrows represent the dependency relationship within a single time slice, where are the dashed arrows denote the temporal evolution of the model between neighboring time slices. $P_{ij}(t_i)$ is the probability of desired goal states $D_1 = i$ and $D_2 = j$ at time t_i. Each node may be expanded as a lower level network, which in turn may be expanded further [32].

what Mission context we are operating in. In this paper, we consider that a mission may consist of sub-goals, which are referred to as tasks/threats. A proactive DSS considers the DMs' workload, time pressure, and role in determining and communicating the relevant information for effective mission performance.

We consider an adaptive model-based approach to further enhance the context and critical event representation. The key benefit of selecting an adaptive model-based approach, informed by data, in lieu of a purely data-driven approach is that the former allows the subject matter expert's knowledge to be incorporated and can be effective from the beginning, even with sparse data, often the case in operational settings, while the latter would require enormous amounts of data spanning a huge variety of mission scenarios to be effective. In addition, when learning graphical models with purely operational data, it is only possible to learn their Markov equivalence classes, meaning that all graphical models in this class represent the same set of conditional independencies, and hence are indistinguishable from each other with respect to the data [33]. Adaptive model-based approaches include graphical models such as dependency graphs (digraphs), Petri nets, multi-functional flow graphs, action-goal attainment (AGA) graphs, hidden Markov models (HMMs), coupled HMMs, factorial HMMs and dynamic hierarchical Bayesian networks (DHBNs) are suitable for representing context-based missions [33]–[37]. Since DHBNs subsume the models embodied in the digraph, Petri nets, multi-functional flow graphs, AGA graphs, and all forms of HMMs, they are good for representing context-based missions. DHBNs are hierarchical directed graphs consisting of a set of nodes and a set of directed edges at various levels of the hierarchy. Each node, representing the contextual elements of the MEAT-H architecture (Figure 3), is considered as a (set of) random variable(s) with either a probability mass function (pmf), probability density function (pdf) for continuous variable(s), or modes to indicate specific states. As the dynamic contextual mission model unfolds over time, it is assumed to be discretized into time slices, where each time slice represents a snapshot of the evolving temporal process/context [32], shown in Figure 4. The nodes, representing the MEAT-H elements, in this network have dependencies with each other (which captures the causal-temporal relationships among the MEAT-H elements). The solid arcs are synchronic to portray the dependency relationships in a single time slice, and the dashed edges are diachronic to show the temporal evolution of the model between neighboring time slices. HMMs can be used to represent the temporal evolution of the states of a node at the lowest level of the hierarchy, as well as other dynamics for network changes from one time slice to the next.

The key components of the DHBN model are shown in Figure 4. The desired objectives are denoted by D, which define the overall mission objectives. The critical/important exogenous events (e.g., environmental conditions, unforeseen mission changes), regarded as context, whose occurrence is beyond the control of the DMs, and affect the evolution of the network states, are denoted by B. A denotes a set of actions (or planned COAs) which can be taken by the DMs to influence the state of the mission. Intermediate goals are not necessarily desired end effects, but are useful in *connecting* the actions and events to the desired end goals, denoted by C. Direct influence dependencies between all the contextual elements and their interactions are specified by conditional probability tables (CPTs) or conditional probability distributions (CPDs) in the parlance of BNs. In order to construct a complex DHBN, we need a generic building block, known as "idioms", which serves as a pattern repository for linking the contextual elements and for displaying the interdependencies of the Bayesian network [38]. An idiom can be of many types (e.g., synthesis, cause-consequence, measurement, induction and reconciliation). Using these idioms, we can model a variety of realistic mission scenarios. For example, a synthesis idiom can be used to represent logical relationships, such as OR, AND, XOR, among multiple objectives under different mission contexts. This enables us to depict the potential context-driven strategies, as well as the dynamic status of actions and objectives. The use of OR nodes delineates context-driven alternative COAs in realizing the DM's intent. The structure and the parameters of the DHBN are instantiated and updated based on operational data.

Learning the network structure enables us to discover new context and to obtain the joint probability of the projected context. However, structure learning is NP-hard, meaning that the computational complexity increases exponentially with the size of the network. Some approximate structure learning methods include reinforcement learning, linear programming relaxation [39], ordering-based search [40], and structural Expectation Maximization [34]. Thus, new data can be incorporated by propagating its effects, as suggested in [34], or by sequentially updating the networks, as in [41]. DHBN-based context representation has the following characteristics: i) finite, but a very large number of states; ii) multi-stage representations, where the decisions are made at the beginning of each stage; iii) stochastic effects generated from action execution; iv) function execution in a particular state results in one of a *number* of possible states with associated probability values due to unforeseen external events (e.g., enemy actions, weather, terrain); and v) Complete or partial observation of the true state of the operational environment at any stage. DHBNs can also be converted into Markov Decision Process (MDP) or Partially Observable Markov Decision Process (POMDP), as needed [42].

B. CONTEXT INFERENCE

The context inference problem in DHBNs is to determine the most likely evolution of nodal states over time based on operational data (evidence). Evidence can be hard (observation of the state of a node) or soft (e.g., pmf associated with the states of a node or sufficient statistics associated with a node) [43]. In the context of counter-smuggling operations, examples of hard evidence include detections, interdictions, asset availability, and current weather information, while soft evidence includes probability of activity (POA) surfaces, flow surfaces based on historical drug interdictions and contraband carried [44], [45]. There are two context inference problems here: an external one related to the Mission, Environment, Assets and Threats, and an internal one related to the DMs, viz., inferring their decision making styles, workload, time pressure, and so on [46], [47].

Compared to the extensive set of exact and approximate inference algorithms for static Bayesian networks, inference algorithms for DHBNs are sparse [48]. Although a transition DHBN can, in theory, be expanded over time, the massive number of nodes generated for a large number of time slices makes this approach impractical. In order to infer external context from operational data of dynamic interactions among contextual elements, the near-optimal inference algorithms based on a novel combination of coordinate ascent, Lagrangian relaxation and Viterbi decoding algorithms developed for coupled HMMs [49]–[55] can be extended to DHBNs. These algorithms decompose the inference problem into decoupled sub-problems, one for each node/MEAT-H element, given the current inference at its child and parent nodes; the sub-problems, which can be solved in parallel, are coordinated by updating the Lagrange multipliers and iterated until convergence. Each sub-problem corresponds to finding the optimal node-state sequence, which can be solved using the Viterbi decoding algorithm [52]–[55]. This approach is ideally suited for distributed, asynchronous implementation. Other approximate inference algorithms include Boyen-Koller [56], loopy belief propagation [33]–[35], and Markov chain Monte Carlo methods [33]–[35]. The external context inference enables us to compare model predictions of outcomes with the expected (desired) outcomes; these deviations form the basis for non-normalcy detection.

C. CONTEXT DETECTION AND PREDICTION

The importance of non-normalcy detection lies in the re-planning stages of the decision making process and forms the basis for root cause analysis (i.e., what caused the change in context) and projection of potential paths that the mission can take given the inferred context and the current plan. The main challenge in context detection is to extract useful features from operational data that can be used to detect anomalies and to quickly estimate the severity of context change with respect to the mission impact.

An adaptive model-based approach (e.g., DHBN) provides a natural mechanism for detecting context changes. Given the current plan, there are expected outcomes or performance measures (e.g., expected number of detections in a surveillance task) based on the DHBN model. The adaptive model-based approach uses residuals (deltas) as features for non-normalcy detection, where the residuals are the outcomes of consistency checks between the actual outcomes derived from operational data and the outputs of the DHBN model. The residuals are large in the presence of context changes and small in the presence of routine or anticipated mission deviations. The residual analysis provides a means to detect changes in context, which triggers (i) root cause analysis [46]–[53], (ii) projection of the impact of the current context on mission goals, and (iii) re-planning strategies to proactively explore decision alternatives to exploit opportunities or mitigate the negative consequences of a changed context. Statistical hypothesis testing techniques (such as change detection [57] e.g., generalized likelihood ratio test, cumulative sum test, sequential probability ratio test, etc.) are used to define thresholds to detect context changes. Additionally, data-driven and knowledge-based approaches can be integrated into an adaptive model-based approach for non-normalcy detection. Data-driven change detection and root cause analysis approaches are derived directly from

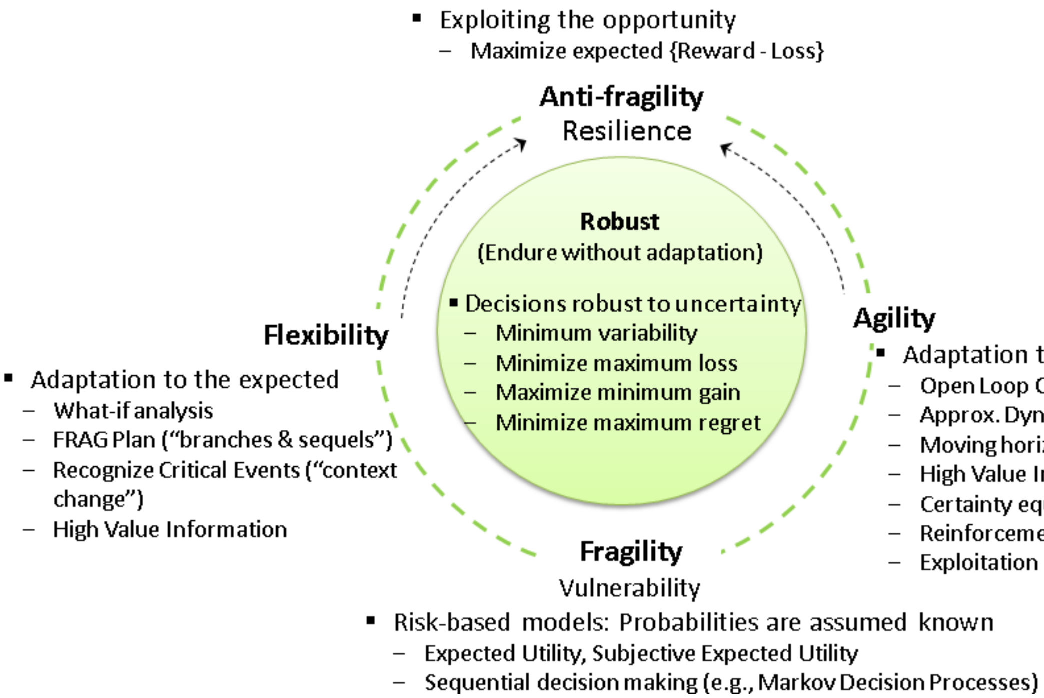

FIGURE 5. Algorithms for uncertainty management.

routinely monitored operating data. The strength of data-driven techniques lies in their ability to transform high-dimensional noisy data into lower dimensional features for detection and diagnostic decisions. Signal analysis methods [58], [59], graphical models [33]–[35], neural networks [35] and multivariate statistical methods [57] are illustrative of data-driven techniques. Knowledge-based approaches are based on qualitative knowledge about the system, where the diagnostic rules are generated using subject matter experts. Thus, model-based/data-driven/knowledge-based approaches can be integrated for robust context change detection.

The context change detection information is further utilized along with the existing mission models (including target and asset models) to predict the probability of achieving the desired objective. As an illustration, in the context of ship routing, we used the predicted weather information to evaluate the amount of wait time at a safe waypoint around a bad weather area. Additionally, the likelihood of reaching the destination can be evaluated depending on the belief in the weather forecasts. Other methods to project context into the future, include Kalman filters, particle filters or DHBN [34], [60], [61].

D. METHODS OF CONTEXT EXTRACTION

Determining the context-driven selection and prioritization of information by quantifying the value of information (VOI) is the key to pre-staging decision-relevant information. When information is successfully valued and extracted, prioritization of contextual information is enabled, thereby promoting mission success. It is therefore important to develop VOI models in order to relay high-value information to DMs. Since each of the MEAT-H elements have associated uncertainty, context becomes inherently uncertain, as well. This exacerbates the challenge of proactive decision making and requires uncertainty management methods in order to efficiently extract and exploit context for decision making. Uncertainty management involves understanding, quantifying and reducing uncertainty for informed decision making to reduce risk (unexpected mission outcomes) and/or maximize reward. Uncertainty management methods, as illustrated in Figure 5, broadly include:

1) *Risk-based methods* which assume that the probability distribution over outcomes is known. Models, such as subjective expected value, subjective expected utility and Markov decision processes that assume the state dynamics and reward structure to be perfectly known, belong to this category. However, these models are vulnerable to sudden or unforeseen changes in context and are therefore fragile in nature [62], [63].
2) *Robust decision making methods* recognize that uncertainty in context exists and thus, seek to manage it by employing conservative decision strategies. These include strategies that minimize variability in the expected risk/reward, minimize maximum risk, maximize minimum reward or minimize maximum regret [64].
3) *Flexible decision making methods* adapt to uncertain context by enumerating or brainstorming potential event sequences a priori, conducting "what-if"

analysis and pre-planning response policies. Typical methods in the context of planning include fragmentary plan (branches and sequels) and conformant plans [65]. These methods are inspired by systems engineering methods such as failure modes, effects and criticality analysis, fault tree analysis, event trees and cause-consequence diagrams, that are used to understand how systems can fail and to mitigate the effects of failures. Thus, flexible decision making methods adapt to expected scenarios by recognizing critical events that signal a context change.

4) *Agile decision making methods* adapt to uncertain context by learning, on-line, an updated model of the decision environment and/or hedging against uncertainty by trading off exploration versus exploitation (also known as dual control or probing and caution in stochastic control literature [66]–[68]). Typical methods in the context of planning include moving horizon planning, certainty equivalence, open-loop optimal feedback, and other approximate dynamic programming (including various forms of reinforcement learning) techniques [66]–[68]. Thus, agile decision making methods adapt to unexpected scenarios by learning the new world model (new context) they are operating in.

5) *Resilient (Anti-fragile) decision making methods* manage uncertainty by adapting to both expected and unexpected events even when information is sparse or lacking. Thus, resilience requires flexibility, robustness and agility. These methods trade-off risk and reward in an uncertain environment by exploiting opportunities. Here, it is important to note that Anti-fragility does not imply that uncertain events will always be experienced positively, rather it means that such systems experience more gains than losses from uncertain events [69].

Thus, a key aspect of uncertainty management is uncertainty reduction by seeking high value information to recognize or learn context changes for informed decision making.

Paucity of information results in poor decisions due to not having enough data pertaining to the context; on the other hand, having too much information will distract and overburden the DM, resulting in poor decision quality [70]. Finding optimal high-value information (HVI) that maximizes the decision quality enables the decision support tool to recommend effective COAs for mission success. A prototype version of this concept is discussed in the context of the counter-piracy problem in [71] and [72], where the impact of POA surface update frequency on the expected decision quality is examined [71]. Similar context-driven analyses in other mission contexts are applicable by considering the uncertainty in POAs themselves (second order uncertainty or ambiguity) and evaluating the sensitivity of HVI with respect to this uncertainty [73]. Our framework employs context-driven selection, extraction, and presentation of information based on its impact on decision quality. This is shown via the arrow feeding into the relevance extractor block in Figure 2, labeled Information Query. This information comprises contextual elements (e.g., the current task is to detect a contact of interest), as well as specific states within the elements of the MEAT-H architecture (e.g., P-3 crew endurance is eight hours). The amount of information and the corresponding level of specificity may vary with the role and cognitive state of the DM (e.g., a targeting board member under time pressure will be presented with only active cases with high payload as opposed to all active and pending cases). In addition to top-down selection of context-relevant HVI, we prioritize the selected information for enhanced context comprehension (Level 2 in Endsley's model of situation awareness [26]). Viable metrics to prioritize the VOI include Bayesian diagnosticity, information gain, and Bayesian optimal experimental design methods [74]–[79], as elaborated on in [80]. Statistical entropy-based computations, including pre-posterior analysis, utility, and Kullback-Leibler Divergence, which also serve as metrics to prioritize information. While these statistical metrics of HVI do not have tangible units, they can serve as a means to rank order multiple pieces of HVI or provide their relative importance. Maximizing decision quality with respect to nodes of a DHBN to be observed, subject to a constraint on the DM's information processing capacity, is similar to the generalized set covering problem used to select optimal sensors for fault detection and isolation [81]. After successfully identifying the context and selecting/prioritizing HVI, the contextualized information can be distributed to the DM(s) via a compact (low bandwidth) format (e.g., Java Script Object Notation, or JSON [82], which can then be relayed to any part, or parts, of the network.

III. PDS APPLICATION SCENARIOS

In this section, we demonstrate the utility of our MEAT-H architecture, HVI models and uncertainty management approaches on three operational problems relevant to the maritime domain: counter-smuggling operations, ship routing, and multiple Unmanned Aircraft System (UAS) mission planning. These maritime missions involve dynamic mission planning (e.g., asset coordination, search path recommendation and asset-task allocation), which requires adaptation to evolving operational context by proactively deducing context-specific COAs for achieving superior mission performance.

A. COUNTER-SMUGGLING OPERATIONS IN THE CONTEXT OF MEAT-H

Counter-smuggling missions involve surveillance operations (to search, detect, track and identify potential threats) and interdiction operations (to intercept, investigate and potentially apprehend suspects). Given the POA surfaces, which integrate meteorological and oceanographic (METOC) and intelligence information (INTEL) to predict where the smugglers may transit, we consider the joint problem of allocating and routing surveillance and interdiction assets to best thwart potential smuggling activities under evolving mission,

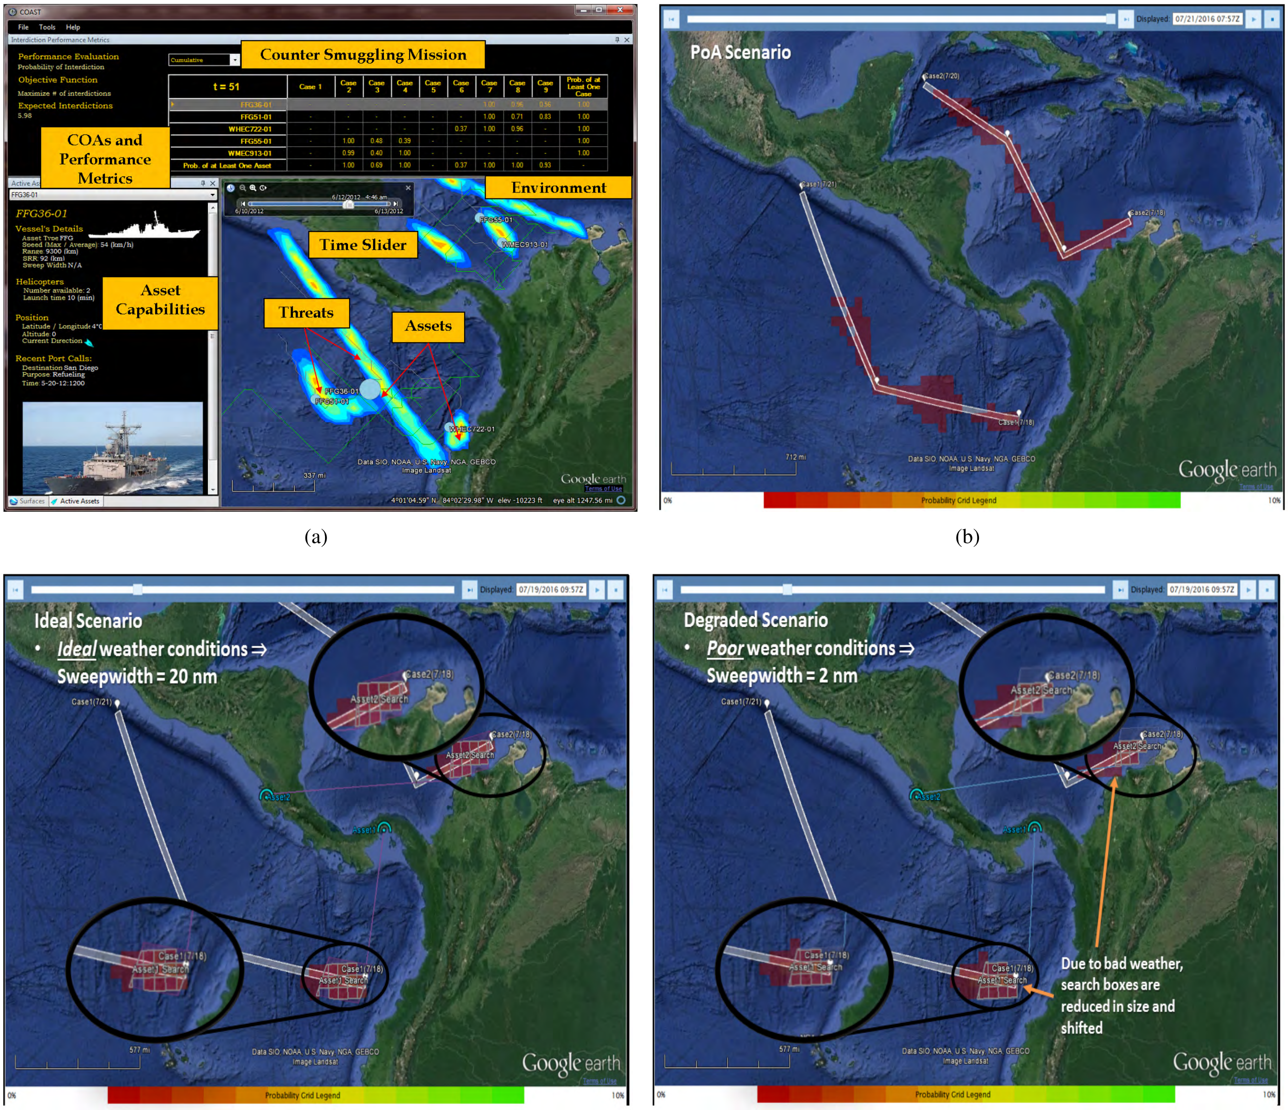

(a) (b) (c) (d)

FIGURE 6. (a) Original cases showing the temporal evolution (as indicated by the color gradient) of the probability of smuggler activity in the East Pacific Ocean and Caribbean Sea. (b) Target corridors (white rectangles) and the uncertainty associated with them in red. (c) Ideal weather scenario with sweep width = 20 nm. (d) Degraded weather conditions impact the performance of the sensors and the sweep width = 2 nm.

environment, asset and threat contexts. Figure 6(a) identifies the MEAT-H elements within the counter-smuggling domain, where the surveillance and interdiction operations are the sub-goals/mission phases to be executed using surveillance and interdiction assets, respectively, such that regions with high POA are maximally covered to prevent smuggling activity, while simultaneously increasing the context awareness. The stochastic counter-smuggling problem (i.e., prone to exogenous events) fits well within the DHBN model as it utilizes the spatio-temporal POA surfaces to search, track and interdict (i.e., intermediate goals) the targets/smugglers within the planning horizon. Since the mission environment is large, the DMs assign surveillance assets to specific search regions and the observations from these assets are processed to characterize the target types, their trajectories and to correlate contacts of interest that have been located with current INTEL. The newly collected information (e.g., INTEL, detections, interdictions, weather data, etc.), serve as a stimuli (which may become a non-normal situation if the values are out of bounds) for context identification. This information is relayed back to the reachback cell in the form of situational reports, using the context protocol (in JSON format). The situational reports are then extracted, processed and aggregated to predict new POA maps (which are considered as HVI) for the next planning interval. The predicted POA surfaces are uncertain and can be prioritized based on weight of contraband to be interdicted or the belief in the target intelligence information. The context-relevant information gathered by the surveillance assets is communicated to the interdiction assets using the context protocol for adapting their COAs to the new context. The coordinated surveillance and interdiction operations may be viewed as a moving horizon stochastic control problem, which is NP-hard (see Appendix A). Our solution approach is employed on open-loop optimal feedback concepts and consists of asymmetric assignments via a Branch and Cut algorithm for the surveillance problem, and an approximate dynamic programming (ADP) algorithm coupled with rollout and Gauss Seidel techniques, to solve the interdiction problem [44], [45], [66], [67], and [83]. These algorithms are embedded in Courses of Action Simulation Tool (COAST), an optimization-based decision support tool for dynamic allocation of surveillance and interdiction assets for counter-smuggling operations. COAST is an optimizer in widget format, integrated with Google Earth, as shown in Figure 6(a).

COAST is proactive in the sense that it automatically incorporates context (with a manual override option) as it has access to operational data, various environment, asset and target models, including the DM's behavior models and inferred context. It allows a flexible targeting strategy by allocating surveillance and interdiction assets based on mission context. As an illustration, Figure 6(b) shows the prior information regarding the target corridors (white rectangles) and the uncertainty (red boxes) related to it, analogous to the possible locations with possible smuggler activity while, taking into account the weather impact, and intelligence information of the approximate time window of departure from a specified port; this is an input to COAST. In the context of calm weather conditions, the sweep width[1] of the surveillance assets is 20 nm, shown in Figure 6(c). However, as the weather worsens (i.e., environmental context change), the sweepwith reduces to 2 nm (change in asset performance models) and this affects the target detection probability of the surveillance asset. Due to the severe weather degradation, the proactive COAST assimilates this change in context and provides modified search boxes (i.e., corrective measure to overcome a non-normal scenario), which are reduced in size and shifted (change in search task), as shown in Figure 6(d). Since the unfavorable weather conditions adversely affect the asset performance, additional surveillance may be required (i.e., change in sub-goals/mission) to concentrate effort on that particular region before routing the interdiction assets to that location. This particular example illustrates how the change in environmental context ripples through the MEAT-H parameters and results in proactive COA recommendations by COAST.

Additional features of proactive COAST allow the DM to specify objectives and the parameters and constraints based on mission context. These objectives include: a) maximizing the number of interdictions/detections over the planning horizon, b) maximizing the contraband interdicted/detected, c) obtaining the highest probability of at least one interdiction/detection, d) maximizing the number of smugglers captured, or e) maximizing the number of unique interdictions/detections. The tool also provides the ability to output the DM-specific m-best solutions for the corresponding objective function, along with valid reasoning, thus allowing the DM to choose the best objective, optimization horizon for that objective, and the different recommendations to output with respect to that objective function. The dynamic coordination between the surveillance and interdiction assets is established via context protocol (in JSON format) for prompt (re)allocation of interdiction assets, once the smuggler has been identified by a surveillance asset. The context-projection and its impact are represented in the form of risk surfaces (contour plots) to the commander and added as a layer to Google Earth to allow for human feedback and solution visualization. These context-specific surfaces form an input to the context-aware design block to find multiple COAs for both interdiction and surveillance problems. In the "what-if" analysis mode, the DM can utilize the tool to explore a variety of mission scenarios by considering different asset types, their capabilities, mission constraints, objectives, and algorithmic options. In scenarios where the mission environment is disrupted, the POA surfaces may not be available or not up-to-date. In such cases, historic (or "flow") POA surfaces are used as input for the asset allocation and scheduling process. Additionally, sensitivity analysis is performed to quantify the value of the POA surfaces and to find the point of diminishing returns. The sensitivity analysis with respect to each of the MEAT-H parameters provides the gradient information on the solution quality, thereby presenting the DMs with cues on how to improve the decision making process. Figure 7 examines the impact of three different types of POA surfaces on the expected decision quality. As shown in Figure 7, providing POA surfaces in the form of

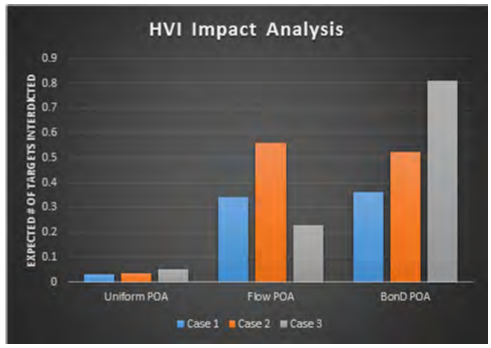

FIGURE 7. Impact analysis of how information types affect mission performance. In this figure, cases refer to targets within a realistic counter-smuggling scenario. In the Uniform POA scenario, representing when little to no information is available, we assumed a uniform distribution of targets across the area of interest. The Flow POA scenario represented when historic information was available, e.g., typical routes traversed by drug smugglers, and resulted in higher quality decision making. Decision quality (measured in terms of the expected number of targets interdicted) was maximized in the third case wherein complete information was available with regards to the POA surfaces (spatial and temporal information providing insight as to *where* and *when* a smuggler will be). Hence, BonD POA surfaces are context-relevant HVI.

[1]Sweep width of any sensor is the width a definite range sensor would have to sweep in order to detect the same number of objects per unit time in a uniform distribution of search objects. It is used to evaluate the probability of detection.

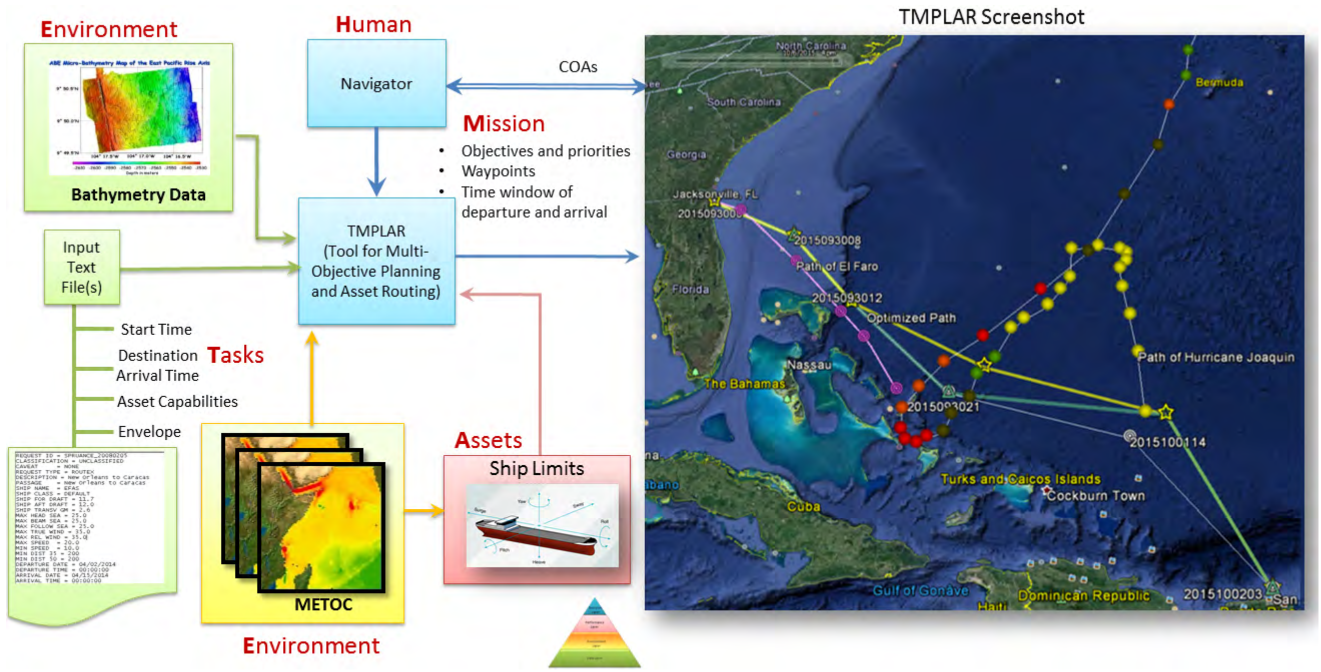

FIGURE 8. Flow diagram of TMPLAR (Tool for Multi-objective PLanning and Asset Routing) which mimics the four layers (Data, Environment, Performance and Decision) of the BonD framework. The TMPLAR screenshot shows hurricane Joaquin (in colored dots where the color represents the strength of the hurricane. TMPLAR screenshot also shows four different routes : a) actual route of El Faro (in purple); b) TMPLAR with original forecast without waiting (in white); c) TMPLAR with original forecast with waiting (in green) and d) TMPLAR with updated forecast with waiting (in yellow) [84].

historical routes of smugglers provided 884% improvement in decision quality (as measured by expected number of targets interdicted) in comparison to a uniform (uninformative) POA surface, while the POA specific to the cases (termed BonD POA) provided 50% further improvement in relation to the historical routes and 1445% improvement overall. Thus, case-based high-valued contextual information in the form of POAs can substantially improve mission success and also facilitate in determining how often to re-plan based on the improvement in decision quality.

B. NAVIGATION IN UNCERTAINTY IN THE CONTEXT OF MEAT

Navigation in uncertainty involves efficient routing of ships in areas impacted by sudden changes in the ocean environment, such as obstacles, weather, fast adversarial boats, etc. Given a graph (e.g., grid maps similar to POA surfaces), a departure point, and a destination point (and possibly a number of waypoints), the objective is to find the shortest path with the lowest cost, where cost may be with respect to a variety of factors (e.g., fuel, fulfillment of training requirements, expected sensor degradation, distance, ship's life). A brief formulation of the multi-objective ship routing problem is provided in Appendix B. In particular, navigation of ships in uncertain environments fits well within our proposed MEAT-H architecture (shown in Figure 8), as this mission involves a number of contextual elements, such as different environmental conditions (bathymetry data, ensemble forecasts with varying spatial-temporal uncertainty over time), multiple objectives, changes in mission goals en route (e.g., training requirements, humanitarian aid) and asset capabilities (ship limits). The MEAT-H elements form an input to the Tool for Multi-objective Planning and Asset Routing (TMPLAR) software [84], which facilitates proactive scheduling of ships (via re-planning or waiting at waypoints to ensure safety) under varying weather and mission contexts and conveys this information to the route planner unobtrusively.

Proactive TMPLAR specifically addresses the following challenges in ship routing under weather uncertainty: 1) Time-dependent costs and asset speed (capable of variable speeds); 2) Time windows of arrival and departure at each waypoint; 3) Uncertain costs and times (due to uncertainty in periodically updated forecasts i.e., predicted METOC); and 4) Costs and times that are functions of the METOC (environmental impacts on ships capabilities, etc.). As an application of proactive decision support, [84], [86] applied TMPLAR to a recent event where a cargo ship (the El Faro), carrying 33 crew members, vanished off the coast of the Bahamas, in Hurricane Joaquin. Figure 8 shows the TMPLAR screenshot with four different routes: a) actual route of El Faro (in purple); b) TMPLAR with original forecast without waiting (in white); c) TMPLAR with original forecast with waiting (in green) and d) TMPLAR with updated forecast with waiting (in yellow). All paths embark at the same time and take nearly the same path for the first 12 hours because the original forecast severely underestimated the strength of Hurricane Joaquin incurring more risk to the cargo ship.

FIGURE 9. Supervisory Control Operations User Testbed (SCOUT): Left screen - shows the position of the UASs, targets in Google Maps, and information about the UASs and targets; Right screen - shows the sensor feed, UAS information (speed, altitude), and two chat windows for INTEL updates and commands [85].

The path El Faro took and the one recommended by TMPLAR without waiting functionality ignore anomalous weather conditions (environment context) and both head into Hurricane Joaquin's path. TMPLAR with waiting using the *original* forecast also succumbs to the hurricane's power because of bad initial forecast. However, proactive TMPLAR with updated forecast waits until the hurricane passes and recommends a route that safely guides the ship to its destination, which could have saved the lives of 33 crew members. This particular application demonstrates when a particular predicted context change (here, weather forecast) becomes significant enough to be incorporated into the decision making process. The routes that TMPLAR found were not only safe for the ship (involving waiting at waypoints and varying the ship's speed in order to avoid the hazardous conditions), but also fuel efficient. Additionally, in scenarios where there is a conflict (e.g., pop-up obstacle or a restricted zone), proactive TMPLAR automatically generates an alternative route, while minimizing the time of route and adhering to the concept of operations.

C. DYNAMIC UAS MISSION PLANNING IN THE CONTEXT OF MEAT-H

In this section, PDS in the cognitive context of human operators is illustrated via Supervisory Control Operations User Testbed (SCOUT), an experimental paradigm developed by the Naval Research Laboratory-Washington DC. SCOUT was designed for the purpose of exploring UAS operator performance in a single operator multiple UAS environment, but assumes some advances in automation necessary to conduct supervisory control operations involving multiple heterogeneous systems. One such advancement includes updating plans based on cognitive context in order to provide COA recommendations to the operators in a proactive and unobtrusive manner. Current unmanned vehicle operations are characterized by teams of operators with highly specialized roles, where the task demands on each operator are independent and highly variable resulting in sub-optimal tasking, mission performance, and mishaps. Maritime operations envision multiple ground-based, aerial, surface and subsurface unmanned vehicles to be simultaneously controlled by a team of operators. In order to achieve this goal, it is crucial to anticipate the future task needs and necessary resource requirements (including the human operators) based on changing environmental conditions. Anticipating future resource requirements based on context changes provides time-critical information on emerging tasks and necessary changes to the order in which the tasks are dynamically assigned and executed, thereby allowing sufficient time for the operators to make appropriate decisions and, in turn, reducing their workload. The key research questions here include: How to determine when the operators are overloaded? And, if they are, what corrective measures (e.g., customized COA recommendations) need to be taken to reduce their cognitive workload?

In order to exploit the PDS framework proposed in this paper, we identify the MEAT-H elements within the SCOUT framework, as shown in Figure 9, and determine the cognitive context, by analyzing the operator's physiological behavior to understand, characterize and predict the cognitive difficulty experienced by them under varying task load levels (i.e., easy, medium and hard workload). The use of psychophysiological measurements (pupillary data, gaze data) as indices of cognitive workload can be considered distinctly superior to other techniques because they can be gathered continuously and coupled with behavioral measures (e.g., risk averse, risk seeker and risk neutral operators) to obtain more information as compared to when using behavioral measures alone. The psychophysiological measurements facilitates in understanding the cognitive context of the human operators and allows proactive adaptation to the changing context. Given the eye-tracking data from NRL's SCOUT, a flexible simulation environment that represents the tasks that a future UAS operator would engage in while controlling multiple UASs, [87], [88] utilized statistical machine learning and classification techniques for cognitive context determination. According to [87] and [88], average pupil dilation, variability of pupil dilation and the power spectral density of pupil dilation are all correlated to cognitive context. By combining several types of features, such as pupil dilation features and gaze metrics, the robustness and accuracy of cognitive context detection can be improved. Additionally, [87], [89] discuss several approaches that can be employed to combine, for example, a series of operator's gaze durations accumulated over a certain period of time (such as mean, median, maximum, standard deviation, etc.) for multiple areas of interest on a screen. In cases where an operator interacts with multiple screens, fusing gaze duration statistics from multiple screens is done via principal component analysis (PCA). PCA is a statistical tool to obtain a smaller number of uncorrelated variables (called principal component scores) that distinguish the data under different conditions, and thus these scores, when used with machine learning algorithms, improve the chances of detecting high workloads with greater accuracy. The results indicate that high workload classification accuracies are possible (e.g., greater than 91% accuracy in distinguishing between easy/medium workload levels and hard workload levels versus 75% accuracy between easy and medium only; this is desirable for adaptive automation). The classification results have the potential in proactively detecting and providing alerts regarding operator overload in realistic UAS mission scenarios. The alerts can be used for improving situation awareness and for providing context-dependent COA recommendations to enhance the productivity of UAS operators, while decreasing operational delays and human fatigue-related mishaps [87], [88].

The information regarding the operator's cognitive context is further utilized to develop proactive recommendations for dynamic scheduling of UASs in order to assist the UAS operators in efficiently managing their workloads. We consider a decision task where an operator must manage multiple UASs and determine the best routes to send their assets to search for targets of varying reward. Targets have some degree of uncertainty associated with their positions, requiring the UAS to search for targets within circular regions of varying radii. During mission execution, the operators update the UAS parameters and target parameters, based upon real-time intelligence provided via chat messages to re-balance workload among the operators, and (re)schedule operator-to-task assignments depending on their cognitive workload. A brief formulation of the dynamic scheduling problem is provided in Appendix C. The dynamic assignment and routing algorithms provide time-critical decision support to the operator on emerging and changing tasks, thereby increasing the time available for human decision making. In particular, the scheduling algorithms (optimal algorithms exemplified by Branch and bound, heuristic approaches such as path time equalization, pairwise exchange combined with rollout) embedded within the SCOUT provide the following capabilities: a) dynamic allocation of targets to UASs in order to meet the target deadlines; b) plan updates based on context

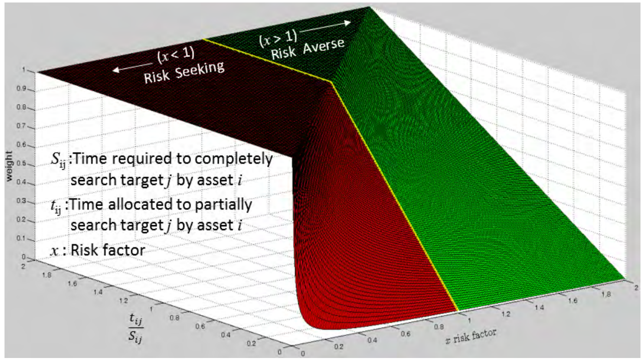

FIGURE 10. Planning under risk. Risk seeker: The weight function rises rapidly with increasing t_{ij}/S_{ij}, i.e. the player assigns a higher value to a target even if the player is able to search the target partially. Risk averse: The growth in weight function with t_{ij}/S_{ij} is much more gradual. Hence a risk-averse player assigns more time to search than required to do 100% search. The yellow line indicates the risk neutral [90].

TABLE 1. Expected utility of targets based on operator's risk propensity.

Target	Reward	Risk Seeking (x<1) (x=0.6)		Risk Neutral(x=1) (x=1)		Risk Averse(x<1) (x=1.4)	
		Area Searched	Expected Utility	Area Searched	Expected Utility	Area Searched	Expected Utility
Kilo	1500	62%	1500	100%	1500	100%	1500
Charlie	1000	100%	1000	89%	890	100%	1000
Bravo	750	88%	750	88%	660	88%	471
Alpha	500	69%	500	29%	148	–	–

changes (e.g., new target, updated information); and c) plan adaptation to operator's risk propensity (viz., risk seeking, risk neutral, risk averse). The proactive SCOUT user interface displays two windows to the operator, as shown in Figure 9. The left screen shows the positions of targets, UASs, and restricted operating zones on the Google map. It has the target information box, which provides the operator with information on the minimum time it takes each UAS to reach a target and the latest time to leave the target after searching. It also has UAS route builder boxes, where an operator assigns targets to UAS(s). The right screen shows the sensor feed, and speed of each UAS. It includes Intelligence and Command chat boxes, which provide the updates on target position, uncertainty radius, and UAS speed during mission execution. UAS operators would benefit greatly from proactive SCOUT to support rapid planning and re-planning, given the likelihood for sub-optimal decisions while handling high workloads. Variables such as target priorities, environmental factors, intelligence uncertainty, etc., have been accounted for in the planning tool. The proactive SCOUT allows the DM to choose the m-best solutions to be displayed on the screen, which provides a means to rank order the performance of UAS operators.

In order to incorporate risk into the decision making process, proactive SCOUT provides customized COA recommendations based on the operator's risk propensity. Risk propensity of an operator (denoted as x in Figure 10) can be broadly classified into three categories: a) risk averse, where targets are devalued if they cannot be completely searched; b) risk neutral, where targets are valued based on the percentage of the corresponding uncertainty region that can be searched; and c) risk seeking, wherein targets with

TABLE 2. Proactive decision support process, algorithms and testbeds.

PDS Operation	Counter Smuggling Operation	Navigation in Uncertainty	Dynamic UAS Mission Planning
Context Representation	**M** Counter-smuggling (surveillance and interdiction operations) **E** East Pacific/ Caribbean sea **A** P-3, E-2C **T** POA Surfaces (Cases, Flow) **H** Decision Makers	**M** Multi-objective routing of assets under uncertainty **E** Area of Interest **A** Ships/ Helicopters/UAS **T** Uncertain weather **H** Decision Makers	**M** Dynamic UAS Scheduling with operator workload balance **E** Area of Interest **A** UASs **T** Pop-up targets **H** Decision Makers
Context Inference for Non-normalcy Detection	• Bad weather conditions • Degraded asset performance • Detection events • Interdiction events	• Uncertain weather conditions • Obstacles • Ship limit exceeds the threshold values	• Pop-up targets • UAS limits • Detection events • Target reward/ uncertainty • Cognitive overload of DMs
Context Prediction	• Predicted METOC information and target behavior are embedded within the updated POA surfaces	• METOC information • Forecast realizations	• DM's workload • COA recommendations based on DM's workload
Context Extraction via Uncertainty Management	• Flexibility to the expected • Agility to the unexpected	• Decisions robust to uncertainty • Exploiting the opportunity	• Exploiting the opportunity • Risk based models • Decisions robust to uncertainty
Algorithms for proactive COA recommendations	• DHBN Modeling • Branch and cut algorithm for surveillance asset allocation • ADP combined with Gauss Seidel and Rollout algorithm for interdiction asset allocation • Sensitivity analysis via HVI impact on mission performance	• Martins' algorithm for context-driven multi-objective optimization • Decision regarding wait time and departure time window from a waypoint while considering predicted weather information	• Neural network algorithms for identifying operator's cognitive workload • Branch and bound algorithm for dynamic UAS scheduling • Heuristic algorithms (path time equalization and pairwise exchange)
Context Visualization / Proactive Decision Support System	• Proactive COAST	• Proactive TMPLAR	• Proactive SCOUT

high rewards (and low probability of complete search) are prioritized over targets with low rewards (and high probability of complete search). A utility theory-based approach is used to model the operator's risk propensity behavior via a weight function, as shown in Figure 10, where the red and green colored regions of the plot correspond to the risk-seeking and the risk-averse behavior, respectively. The risk-neutral behavior is denoted by the yellow line (i.e., $x = 1$). In case of a risk-seeking operator, the value of the weighing function rises rapidly with the increase in t_{ij}/S_{ij}, whereas this rise is very gradual for a risk-averse operator. This implies that a risk-averse operator spends more time than a risk-seeker to obtain the same reward. Table 1 illustrates this concept, by considering a scenario where the target Bravo has a maximum achievable reward of 750 units and only 88% of its area can be searched within the stipulated deadline. Based on the risk-propensity model, a risk-seeking operator ($x = 0.6$) would expect the maximum achievable utility of 750 units, where as a risk-averse operator ($x = 1.4$) devalues this target to an expected utility of 471 units (=0.88 × 750/1.4). The risk-neutral operator assigns it a utility of 660 units (=0.88 × 750) [90]. Therefore, depending on the operator's past behavior (risk averse, seeking or neutral), the COA recommendations are customized accordingly. Thus, proactive SCOUT identifies and anticipates the cognitive context of human operators engaged in a realistic operation and presents context-relevant COAs in an unobtrusive manner to assure that the operators are attending to the right task at the right time such that task demands do not exceed operator capabilities in a multi-mission environment.

IV. CONCLUSION AND FUTURE WORK

In this paper, we developed a generalized context-driven PDS which is applicable across multiple maritime missions, as summarized in Table II. The PDS framework and the algorithms discussed in the paper facilitate the identification, extraction and quantification of the context-specific information, along with proactive presentation and dissemination of this information to DMs for faster, context-specific and anticipatory decision making. Additionally, it allows the DMs to promptly understand and envision the current and projected mission context, while allowing them ample time to make appropriate decisions by taking into consideration the concomitant uncertainties, and unknown risks stemming from the specific context via networking, collaboration, distributed execution, and resource sharing within the mission environment. In particular, the PDS framework contributes to the maritime mission planning processes by: 1) increasing the efficiency of asset utilization at the tactical and the operational levels; 2) dynamically invoking pre-planned COAs as a function of emerging events, while automatically adapting plans to unfolding events, and rapid re-planning in the event of unforeseen context evolution; 3) timely, relevant and validated decision support for operational level commanders by moving the right information/knowledge to the right people at the right time to enable proactive information processing; and, 4) reducing the operator overload via dynamic tasking and scheduling in a decentralized environment. Our future research includes extending the PDS framework for theater mission planning in littoral environments via incorporation of multistage pursuit-evasion games (e.g., Stackelberg security games) that utilize POA surfaces as input and provide proactive COAs in disrupted and low-bandwidth environments. We plan to explore other viable methods, such as deep reinforcement opponent learning, to infer the adversary's intent. Additionally, we plan to develop parallel algorithms in order to infer context and notify DMs if mission performance parameters breach certain thresholds, while the mission is in progress.

APPENDIX A

The mathematical formulation for the counter-piracy operation (also extensible to the counter-smuggling operation) can be formulated as a stochastic control problem involving surveillance and interdiction sub-problems [71], [72]. The nomenclature is detailed in Table III.

A. SURVEILLANCE SUBPROBLEM

$$\max_{\{A_s(k), s\in S_k\}} \sum_{k=1}^{K} \gamma^{(k-1)} \sum_{s\in S_k} PD_s(A_s(k), k)$$
$$\text{s.t. } A_i(k) \cap A_j(k) = \emptyset, \quad (i \neq j) \in S_k$$
$$A_i(k) \text{ has rectangular shape}, \quad \forall i \in S_k \tag{1}$$

In the above objective function, the probability of detection, denoted PD_s, is defined as

$$PD_s(A_s(k), k) = \sum_{g\in A_s(k)} PP(g,k) \times \left(1 - \exp\left(\frac{-w_s(k) * v_s(k) * \tau_s(k)}{a_c |A_s(k)|}\right)\right) \tag{2}$$

TABLE 3. Summary of notation for surveillance and interdiction problems.

g	Cell index contained within G
k	Time epoch index
i, s, j	Asset index
K	Final period to be planned for (end of horizon)
I_k	Set of available interdiction assets during time epoch k
S_k	Set of available surveillance assets during time epoch k
$x_i(k)$	Location of asset i during time epoch k
v_i	Maximum speed for interdiction asset i
v_i^h	Maximum speed for helicopter h on-board asset i
a_c	Area of cell
t_i^h	Launch delay time for helicopter h
τ	Time to reach the target
$\tau_s(k)$	Time spent by asset s in the search area
R_i	Range of reachable cells for asset i
$\rho(i)$	Surface radar range of asset i
$w_s(k)$	Sweep width of asset s at time k
$v_s(k)$	Maximum speed for surveillance asset s at time k
$A_s(k)$	Area assigned to asset s at time k (in number of cells)
$r(i,\tau)$	Distance that can be covered during τ in a coordinated effort between interdiction asset i and helicopter h
PP	Probability of pirate presence
PI	Probability of interdiction
PD	Probability of detection
γ	Discount factor (monotonically non-increasing)
$PoA(g,k)$	Probability activity in cell g during time epoch k
$dist(x_i(k), g)$	Euclidean distance from cell g to the location of asset i, $x_i(k)$

B. INTERDICTION SUBPROBLEM

$$\max_{\mathbf{U}} \sum_{k=1}^{K} \gamma^{(k-1)} \sum_{g\in G} PoA(g,k) \bullet \left[1 - \prod_{i\in I_k} (1 - PI_i(x_i(k), g)\, PD(g,k))\right]$$
$$s.t.\ x_i(k+1) \in R_i[x_i(k)] \tag{3}$$

where

$$\mathbf{U} = \{x_i(k)\ \forall k = 1, 2, .., K,\ \forall i \in I_k\}$$

$$PD(g,k) = \begin{cases} 1, & dist(x_i(k), g) \le \rho(i) \\ 0, & \text{otherwise} \end{cases} \tag{4}$$

$$PI_i(x_i(k), g) = \begin{cases} \dfrac{2r(i,\tau)}{dist(x_i(k), g)}, & r(i,\tau) < \dfrac{dist(x_i(k), g)}{2} \\ 1, & r(i,\tau) \ge \dfrac{dist(x_i(k), g)}{2} \end{cases} \tag{5}$$

here

$$r(i,\tau) = \begin{cases} v_i\tau, & \tau \le t_i^h \\ v_i t_i^h + v_i^h(\tau - t_i^h), & \tau > t_i^h \end{cases}$$

$R_i[x_i(k)]$ represents the set of reachable cells by asset i within time period $k+1$ given the current location at time k.

TABLE 4. Summary of notation for multi-objective ship routing.

Discrete Variables:	k	Time epoch
	i	Specific objective
	d	Total number of objectives
State Variables:	$\underline{x} = (G, k)$, where G is navigational coordinates	
Control Variables:	$\underline{u} = (\psi, P)$, where ψ denotes a set of allowable transitions between present states and the states in the next time epoch	
	P	Discretized power output
Motion Dynamics:	$\underline{M}$	Generalized ship motion

TABLE 5. Summary of notation for dynamic scheduling of UASs.

$\mathbb{T}$	Set of targets
$\mathbb{U}$	Team of UASs
i	Asset index
j, k	Target index or dummy index
m	Total number of UASs
ψ_i	Dummy index (Depot)
s_{ik}	Time spent by UAS u_i in searching for target t_k within the target's uncertainty region
ψ_i	Dummy index (depot)
ξ_{ijk}	Binary (0-1) decision variable
τ_k	Arrival time of assigned UAS at target t_k
D_k	Deadline for target t_k
MT	Mission time
E	Cumulative utility
$E_{ik}(\tau_k)$	Utility obtained from target t_k
p_{ijk}	Time taken by UAS u_i to travel from target t_j to target t_k

APPENDIX B

Table IV lists the notation summary for the multi-objective ship routing problem, which is formulated formally in [84].

$$\begin{aligned} &\min\ J(\underline{x}) \\ &s.t.\ \underline{x}_k \in \underline{X}_a(k) \quad \text{Within the predefined grid system} \\ &\quad \underline{u}_k \in \underline{U}_a(k) \quad \text{Allowable transitions between grid points} \\ &\quad \underline{M}_k \in \underline{M}_a(k) \quad \text{Allowable motion} \end{aligned} \tag{6}$$

where $J(\underline{x}) = [J^1(\underline{x}), \ldots, J^i(\underline{x}), \ldots, J^d(\underline{x})]$; and $J^i(\underline{x}) = \sum_{k=1}^{N-1} g_k^i(\underline{x}_k, \underline{u}_k, \underline{M}_k) + g_N^i(\underline{x}_N)$ is the total cost for an N-stage path planning. Here, g_k^i represents the cost (with respect to objective i) of being at a node, while g_N^i denotes the terminal cost corresponding to objective i.

APPENDIX C

The notation summary for dynamic scheduling of UASs is listed in Table V. The dynamic scheduling of UASs is formulated as a bi-objective problem of maximizing the cumulative utility, while simultaneously minimizing the mission time (makespan). The constraints ensure that not every target is assigned to a UAS, while each UAS is assigned only once and the number of UASs assigned must not exceed the number of available UASs. The arrival time and deadline constraints are also taken into account [90].

$$\begin{aligned} &\max\ J : \max \sum_{i\in\mathbb{U}} \sum_{j\in\psi_i\cup\mathbb{T}} \sum_{k\in\mathbb{T}} E_{ik}(\tau_k)\, \xi_{ijk} \\ &\min\ MT : \min \max_{i\in\mathbb{U},k\in\mathbb{T}} (\tau_k + s_{ik}) \\ &\text{s.t.} \sum_{i\in\mathbb{U}} \sum_{j\in\psi_i\cup\mathbb{T}} \xi_{ijk} \le 1 \\ &\quad \sum_{k\in\mathbb{T}} \xi_{i\psi_i k} \le 1 \\ &\quad \sum_{i\in\mathbb{U}} \sum_{k\in\mathbb{T}} \xi_{i\psi_i k} \le m \\ &\quad \tau_k = \sum_{i\in\mathbb{U}} \sum_{j\in\psi_i\cup\mathbb{T}} \left(\tau_j + s_{ij}(\tau_j) + p_{ijk}\right) \times \xi_{ijk} \\ &\quad 0 \le \tau_k \le D_k \\ &\quad \tau_k - \tau_j > 0, \quad \forall \xi_{ijk} = 1 \\ &\quad \tau_{\psi_i} = 0, \quad \forall i \end{aligned} \tag{7}$$

where

$$\xi_{ijk} = \begin{cases} 1, & \text{if UAS } u_i \text{ is assigned to target } t_k \text{ immediately after target } t_j \text{ or UAS's initial position} \\ 0, & \text{otherwise} \end{cases}$$

ACKNOWLEDGMENT

The authors would like to thank NRL-Monterey and Joint Interagency Task Force-South (JIATF-S) for their valuable contributions in designing and detailing the drug trafficking scenarios and the concept of operations embedded within our decision support software tool. The authors would also like to thank NRL-DC for providing the SCOUT software as a testbed for validating our algorithms. The authors are thankful to Dr. Jeffrey G. Morrison and the Command Decision Making program at the U.S. Office of Naval Research.

REFERENCES

[1] C. J. McCarthy and U.S. Air Force, "Anti-access/area denial: The evolution of modern warfare," Joint Military Oper. Dept., U.S Naval War College, Tech. Rep., 2010, p. 73. [Online]. Available: http://www.dtic.mil/dtic/tr/fulltext/u2/a525078.pdf

[2] K. Neville, B. Blickensderfer, J. Archer, K. Kaste, and S. P. Luxion, "A cognitive work analysis to identify human-machine interface design challenges unique to uninhabited aircraft systems," in *Proc. Human Factors Ergonomics Soc. Annu. Meeting*, vol. 56. 2012, pp. 418–422.

[3] UAV crash blamed on distraction of battle, *Air Force Times*, Vienna, VA, USA, 2010.

[4] X. Han, H. Bui, S. Mandal, K. R. Pattipati, and D. L. Kleinman, "Optimization-based decision support software for a team-in-the-loop experiment: Asset package selection and planning," *IEEE Trans. Syst., Man, Cybern., Syst.*, vol. 43, no. 2, pp. 237–251, Mar. 2013.

[5] M. Mishra, W. An, X. Han, D. Sidoti, D. F. M. Ayala, and K. R. Pattipati, "Decision support software for anti-submarine warfare mission planning within a dynamic environmental context," in *Proc. IEEE Int. Conf. Syst., Man*, Oct. 2014, pp. 3390–3393.

[6] CNMOC. *Battlespace on Demand Framework*, accessed on 2010. [Online]. Available: http://www.dtcenter.org/events/workshops10/nuopc/presentations/payne.pdf

[7] P. Kime, "Navy should bolster crisis planning for theater ASW," *Sea Power*, vol. 46, no. 9, p. 12, 2003.

[8] A. Smirnov, "Context-driven decision making in network-centric operations: Agent-based intelligent support," Russian Acad. Sci. Saint Petersburg Inst. Inf. Autom., Saint Petersburg, Russia, Tech. Rep., 2006.

[9] B. N. Schilit and M. M. Theimer, "Disseminating active map information to mobile hosts," *IEEE Netw.*, vol. 8, no. 5, pp. 22–32, Sep. 1994.

[10] P. J. Brown, "The stick-e document: A framework for creating context-aware applications," *Electron. Publishing-Chichester*, vol. 8, nos. 2–3, pp. 259–272, 1995.

[11] P. J. Brown, J. D. Bovey, and X. Chen, "Context-aware applications: From the laboratory to the marketplace," *IEEE Pers. Commun.*, vol. 4, no. 5, pp. 58–64, Oct. 1997.

[12] A. Ward, A. Jones, and A. Hopper, "A new location technique for the active office," *IEEE Pers. Commun.*, vol. 4, no. 5, pp. 42–47, Oct. 1997.

[13] P. Korpipää, J. Mäntyjärvi, J. Kela, H. Keränen, and E.-J. Malm, "Managing context information in mobile devices," *IEEE Pervas. Comput.*, vol. 2, no. 3, pp. 42–51, Jul./Sep. 2003.

[14] T. Rodden, K. Chervest, K. Davies, and A. Dix, "Exploiting context in HCI design for mobile systems," in *Proc. Workshop Human Comput. Interact. Mobile Devices*, Glasgow, Scotland, 1998, pp. 21–22.

[15] B. Schilit, N. Adams, and R. Want, "Context-aware computing applications," in *Proc. 1st Workshop Mobile Comput. Syst. Appl.*, Dec. 1994, pp. 85–90.

[16] G. D. Abowd, A. K. Dey, P. J. Brown, N. Davies, M. Smith, and P. Steggles, "Towards a better understanding of context and context-awareness," in *Proc. Int. Symp. Handheld Ubiquitous Comput.*, 1999, pp. 304–307.

[17] A. K. Dey, G. D. Abowd, and D. Salber, "A conceptual framework and a toolkit for supporting the rapid prototyping of context-aware applications," *Human-Comput. Interact.*, vol. 16, nos. 2–4, pp. 97–166, 2001.

[18] P. Brézillon, "Context in problem solving: A survey," *Knowl. Eng. Rev.*, vol. 14, no. 1, pp. 47–80, 1999.

[19] O. Kwon, S. Choi, and G. Park, "NAMA: A context-aware multi-agent based Web service approach to proactive need identification for personalized reminder systems," *Expert Syst. Appl.*, vol. 29, no. 1, pp. 17–32, Jul. 2005.

[20] O. Kwon, K. Yoo, and E. Suh, "UbiDSS: A proactive intelligent decision support system as an expert system deploying ubiquitous computing technologies," *Expert Syst. Appl.*, vol. 28, no. 1, pp. 149–161, Jan. 2005.

[21] T. R. Gruber, "A translation approach to portable ontology specifications," *Knowl. Acquisition*, vol. 5, no. 2, pp. 199–220, Jun. 1993.

[22] J.-Y. Hong, E.-H. Suh, and S.-J. Kim, "Context-aware systems: A literature review and classification," *Expert Syst. Appl.*, vol. 36, no. 4, pp. 8509–8522, May 2009.

[23] G. Chen and D. Kotz, "A survey of context-aware mobile computing research," Dept. Comput. Sci., Dartmouth College, Hanover, NH, USA, Tech. Rep. TR2000-381, 2000, pp. 68–72.

[24] C. Bettini *et al.*, "A survey of context modelling and reasoning techniques," *Pervas. Mobile Comput.*, vol. 6, no. 2, pp. 161–180, Apr. 2010.

[25] H. A. Simon, *The New Science of Management Decision*. New York, NY, USA: Harper & Row, 1960.

[26] M. R. Endsley, "Toward a theory of situation awareness in dynamic systems," *Human Factors, J. Human Factors Ergonom. Soc.*, vol. 37, no. 1, pp. 32–64, 1995.

[27] N. N. Taleb, *Antifragile: Things That Gain From Disorder*, vol. 3. Westminster, MD, USA: Random House Inc., 2012.

[28] G. M. Levchuk, Y. N. Levchuk, J. Luo, K. R. Pattipati, and D. L. Kleinman, "Normative design of organizations. I. Mission planning," *IEEE Trans. Syst., Man, Cybern. A, Syst., Humans*, vol. 32, no. 3, pp. 346–359, May 2002.

[29] G. M. Levchuk, Y. N. Levchuk, J. Luo, K. R. Pattipati, and D. L. Kleinman, "Normative design of organizations. II. Organizational structure," *IEEE Trans. Syst., Man, Cybern. A, Syst., Humans*, vol. 32, no. 3, pp. 360–375, May 2002.

[30] G. M. Levchuk, Y. N. Levchuk, C. Meirina, K. R. Pattipati, and D. L. Kleinman, "Normative design of project-based organizations—Part III: Modeling congruent, robust, and adaptive organizations," *IEEE Trans. Syst., Man, Cybern. A, Syst., Humans*, vol. 34, no. 3, pp. 337–350, May 2004.

[31] N. F. Noy, M. Sintek, S. Decker, M. Crubezy, R. W. Fergerson, and M. A. Musen, "Creating semantic Web contents with protege-2000," *IEEE Intell. Syst.*, vol. 16, no. 2, pp. 60–71, Mar./Apr. 2001.

[32] H. Tu, Y. N. Levchuk, and K. R. Pattipati, "Robust action strategies to induce desired effects," *IEEE Trans. Syst., Man, Cybern. A, Syst., Humans*, vol. 34, no. 5, pp. 664–680, Sep. 2004.

[33] D. Koller and N. Friedman, *Probabilistic Graphical Models: Principles and Techniques*. Cambridge, MA, USA: MIT Press, 2009.

[34] K. P. Murphy, *Machine Learning: A Probabilistic Perspective*. Cambridge, MA, USA: MIT Press, 2012.

[35] C. M. Bishop, *Pattern Recognition and Machine Learning*, vol. 128. Berlin, Germany: Springer, 2006.

[36] S. Deb, K. R. Pattipati, V. Raghavan, M. Shakeri, and R. Shrestha, "Multi-signal flow graphs: A novel approach for system testability analysis and fault diagnosis," *IEEE Aerosp. Electron. Syst. Mag.*, vol. 10, no. 5, pp. 14–25, May 1995.

[37] K. Jensen, "Coloured petri nets," in *Petri Nets: Central Models and Their Properties*. Berlin, Germany: Springer, 1987, pp. 248–299.

[38] N. Fenton and M. Neil, *Risk Assessment and Decision Analysis With Bayesian Networks*. Boca Raton, FL, USA: CRC, 2012.

[39] T. S. Jaakkola *et al.*, "Learning Bayesian network structure using LP relaxations," in *Proc. AISTATS*, 2010, pp. 358–365.

[40] M. Teyssier and D. Koller. (2012). "Ordering-based search: A simple and effective algorithm for learning Bayesian networks." [Online]. Available: https://arxiv.org/abs/1207.1429

[41] N. Friedman and M. Goldszmidt, "Sequential update of Bayesian network structure," in *Proc. 13th Conf. Uncertainty Artif. Intell.*, 1997, pp. 165–174.

[42] C. Meirina, Y. N. Levchuk, G. M. Levchuk, K. R. Pattipati, and D. L. Kleinman, "A markov decision process (MDP) approach to goal attainment," *IEEE Trans. Syst., Man, Cybern. A, Syst., Humans*, vol. 38, no. 1, pp. 116–133, Jan. 2008.

[43] H. Tu, J. Allanach, S. Singh, K. R. Pattipati, and P. Willett, "Information integration via hierarchical and hybrid Bayesian networks," *IEEE Trans. Syst., Man, Cybern. A, Syst., Humans*, vol. 36, no. 1, pp. 19–33, Jan. 2006.

[44] D. Sidoti *et al.*, "Context-aware dynamic asset allocation for maritime interdiction operations—Part I: Approximate dynamic programming approaches," *J. Adv. Inf. Fusion*, to be published.

[45] D. Sidoti *et al.*, "Context-aware dynamic asset allocation for maritime interdiction operations—Part II: Application and sensitivity analysis," *IEEE Trans. Syst., Man, Cybern., Syst.*, to be submitted.

[46] V. C. Trinh and A. J. Gonzalez, "Discovering contexts from observed human performance," *IEEE Trans. Human-Mach. Syst.*, vol. 43, no. 4, pp. 359–370, Jul. 2013.

[47] R. Daly, Q. Shen, and S. Aitken, "Learning Bayesian networks: Approaches and issues," *Knowl. Eng. Rev.*, vol. 26, no. 2, pp. 99–157, 2011.

[48] H. K. G. Fernlund, A. J. Gonzalez, M. Georgiopoulos, and R. F. DeMara, "Learning tactical human behavior through observation of human performance," *IEEE Trans. Syst. Man, Cybern. B, Cybern.*, vol. 36, no. 1, pp. 128–140, Feb. 2006.

[49] M. Shakeri, V. Raghavan, K. R. Pattipati, and A. Patterson-Hine, "Optimal and near-optimal algorithms for multiple fault diagnosis with unreliable tests," in *Proc. Test Technol. Commercialization. Conf. Rec.*, 1996, pp. 473–482.

[50] F. Tu, K. R. Pattipati, S. Deb, and V. N. Malepati, "Multiple fault diagnosis in graph-based systems," *Proc. SPIE*, vol. 4733, pp. 168–179, Jul. 2002. [Online]. Available: http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=887514

[51] F. Yu, F. Tu, H. Tu, and K. R. Pattipati, "A Lagrangian relaxation algorithm for finding the MAP configuration in QMR-DT," *IEEE Trans. Syst., Man, Cybern. A, Syst., Humans*, vol. 37, no. 5, pp. 746–757, Sep. 2007.

[52] S. Singh *et al.*, "Dynamic multiple fault diagnosis: Mathematical formulations and solution techniques," *IEEE Trans. Syst., Man, Cybern. A, Syst., Humans*, vol. 39, no. 1, pp. 160–176, Jan. 2009.

[53] A. Kodali, S. Singh, and K. Pattipati, "Dynamic set-covering for real-time multiple fault diagnosis with delayed test outcomes," *IEEE Trans. Syst., Man, Cybern., Syst.*, vol. 43, no. 3, pp. 547–562, May 2013.

[54] A. Kodali, K. Pattipati, and S. Singh, "A coupled factorial hidden Markov model (CFHMM) for diagnosing coupled faults," in *Proc. IEEE Aerosp. Conf.*, Mar. 2010, pp. 1–11.

[55] S. Zhang, K. R. Pattipati, Z. Hu, X. Wen, and C. Sankavaram, "Dynamic coupled fault diagnosis with propagation and observation delays," *IEEE Trans. Syst., Man, Cybern., Syst.*, vol. 43, no. 6, pp. 1424–1439, Nov. 2013.

[56] X. Boyen and D. Koller, "Approximate learning of dynamic models," in *Proc. 11th Adv. Neural Inf. Process. Syst.*, vol. 11. 1999, p. 396.
[57] M. Basseville, "Detecting changes in signals and systems—A survey," *Automatica*, vol. 24, no. 3, pp. 309–326, May 1988.
[58] M. Azam, K. Pattipati, J. Allanach, S. Poll, and A. Patterson-Hine, "In-flight fault detection and isolation in aircraft flight control systems," in *Proc. IEEE Aerosp. Conf.*, Mar. 2005, pp. 3555–3565.
[59] M. S. Azam, F. Tu, K. R. Pattipati, and R. Karanam, "A dependency model-based approach for identifying and evaluating power quality problems," *IEEE Trans. Power Del.*, vol. 19, no. 3, pp. 1154–1166, Jul. 2004.
[60] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation With Applications to Tracking and Navigation: Theory Algorithms and Software*. New York, NY, USA: Wiley, 2004.
[61] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking," *IEEE Trans. Signal Process.*, vol. 50, no. 2, pp. 174–188, Feb. 2002.
[62] D. Kahneman and A. Tversky, "Prospect theory: An analysis of decision under risk," *Econometrica, J. Econ. Soc.*, vol. 47, no. 2, pp. 263–291, 1979.
[63] S. Barberà, P. Hammond, and C. Seidl, Eds., *Handbook of Utility Theory: Extensions*, vol. 2. Berlin, Germany: Springer, 2004.
[64] E. L. Lawler, *Combinatorial Optimization: Networks and Matroids*. North Chelmsford, MA, USA: Courier Corporation, 2001.
[65] J. Kurien, P. Nayak, and D. E. Smith, "Fragment-based conformant planning," in *Proc. AIPS*, 2002, pp. 153–162.
[66] W. B. Powell, *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, vol. 703. New York, NY, USA: Wiley, 2007.
[67] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, vol. 1. Cambridge, U.K.: MIT Press, 1998, p. 1.
[68] D. P. Bertsekas, *Dynamic Programming and Optimal Control: Approximate Dynamic Programming*, vol. 2. Belmont, MA, USA: Athena Scientific, 2012, p. 4.
[69] N. N. Taleb and R. Douady, "Mathematical definition, mapping, and detection of (anti)fragility," *Quant. Finance*, vol. 13, no. 11, pp. 1677–1689, 2013.
[70] M. S. John, H. S. Smallman, D. I. Manes, B. A. Feher, and J. G. Morrison, "Heuristic automation for decluttering tactical displays," *Human Factors, J. Human Factors Ergonom. Soc.*, vol. 47, no. 3, pp. 509–525, 2005.
[71] D. Sidoti *et al.*, "Evaluating the value of information in the presence of high uncertainty," presented at the 18th Int. Command Control Res. Technol. Symp., C2 Underdeveloped, Degraded, Denied Oper. Environ., Alexandria, VA, USA, Jun. 2013.
[72] W. An *et al.*, "Dynamic asset allocation approaches for counter-piracy operations," in *Proc. 15th Int. Conf. Inf. Fusion*, Jul. 2012, pp. 1284–1291.
[73] H. J. Einhorn and R. M. Hogarth, "Ambiguity and uncertainty in probabilistic inference," *Psychol. Rev.*, vol. 92, no. 4, pp. 433–461, 1985.
[74] J. Baron, *Rationality and Intelligence*. Cambridge, U.K.: Cambridge Univ. Press, 2005.
[75] L. Stein, "Probability and the weighing of evidence," *Brit. J. Soc. Med.*, vol. 4, no. 3, pp. 170–171, 1950.
[76] J. Klayman and Y.-W. Ha, "Confirmation, disconfirmation, and information in hypothesis testing," *Psychol. Rev.*, vol. 94, no. 2, pp. 211–228, 1987.
[77] S. Kullback and R. A. Leibler, "On information and sufficiency," *Ann. Math. Stat.*, vol. 22, no. 1, pp. 79–86, 1951.
[78] D. V. Lindley, "On a measure of the information provided by an experiment," *Ann. Math. Stat.*, vol. 27, no. 4, pp. 986–1005, 1956.
[79] J. D. Nelson, "Finding useful questions: On Bayesian diagnosticity, probability, impact, and information gain," *Psychol. Rev.*, vol. 112, no. 4, p. 979, 2005.
[80] J. D. Nelson, "Towards a rational theory of human information acquisition," in *The Probabilistic Mind: Prospects for Bayesian Cognitive Science*. Oxford, U.K.: Oxford Univ. Press 2008, pp. 143–163.
[81] S. Zhang, K. R. Pattipati, Z. Hu, and X. Wen, "Optimal selection of imperfect tests for fault detection and isolation," *IEEE Trans. Syst., Man, Cybern., Syst.*, vol. 43, no. 6, pp. 1370–1384, Nov. 2013.
[82] T. Bray, "The JavaScript object notation (JSON) data interchange format," Tech. Rep., 2014.
[83] D. P. Bertsekas, J. N. Tsitsiklis, and C. Wu, "Rollout algorithms for combinatorial optimization," *J. Heuristics*, vol. 3, no. 3, pp. 245–262, 1997.
[84] D. Sidoti *et al.*, "A multiobjective path-planning algorithm with time windows for asset routing in a dynamic weather-impacted environment," *IEEE Trans. Syst., Man, Cybern., Syst.*, to be published.
[85] C. Sibley, J. Coyne, and J. Thomas, "Demonstrating the supervisory control operations user testbed (SCOUT)," in *Proc. Human Factors Ergonomics Soc. Annu. Meeting*, vol. 60. 2016, pp. 1324–1328.
[86] D. Sidoti, "Why context and context-driven decision support matters," *IEEE Aerosp. Electron. Syst. Mag.*, vol. 31, no. 11, pp. 50–52, Nov. 2016.
[87] P. Mannaru, B. Balasingam, K. Pattipati, C. Sibley, and J. Coyne, "On the use of hidden Markov models for gaze pattern modeling," *Proc. SPIE*, vol. 9851, p. 98510R, May 2016. [Online]. Available: http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=2523341
[88] P. Mannaru, B. Balasingam, K. Pattipati, C. Sibley, and J. Coyne, "Cognitive context detection using pupillary measurements," *Proc. SPIE*, vol. 9851, p. 98510Q, May 2016. [Online]. Available: http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=2523340
[89] P. Mannaru, B. Balasingam, K. Pattipati, C. Sibley, and J. Coyne, "Cognitive context detection in UAS operators using eye-gaze patterns on computer screens," *Proc. SPIE*, vol. 9851, p. 98510F, May 2016. [Online]. Available: http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=2523329
[90] B. K. Nadella *et al.*, "Proactive decision support for dynamic assignment and routing of unmanned aerial systems," in *Proc. IEEE Aerosp. Conf.*, Mar. 2016, pp. 1–11.

MANISHA MISHRA received the B.S. degree in electrical engineering from Punjab Technical University, India, in 2004, and the M.S. degree in electrical and computer engineering from the University of Hawaii at Manoa, Honolulu, HI, USA, in 2008. She is currently pursuing the Ph.D. degree with the Department of Electrical and Computer Engineering, University of Connecticut, Storrs, CT, USA. She was a Software Engineer with Infosys Technologies, Bengaluru, India, in 2005. She is a Post-Graduate Researcher with the U.S. Army Engineering Research and Development Center-Cold Region Research Labs, Hanover, NH, USA. Her current research interests include modeling dynamic and uncertain environments for asset allocation and multiobjective path planning, context aware decision support systems, risk analysis, system diagnosis and prognosis, and optimization-based techniques for mission planning and coordination.

DAVID SIDOTI received the B.S. and M.S. degrees in electrical and computer engineering from the University of Connecticut, Storrs, CT, USA, in 2011 and 2016, respectively, where he is currently pursuing the Ph.D. degree with the Department of Electrical and Computer Engineering, under the advisement of Dr. K. R. Pattipati. His research in dynamic resource management has resulted in numerous transitions to the real world (Naval Research Laboratory-Monterey and Joint Interagency Task Force-South). He has authored or co-authored over 30 research articles in peer-reviewed journals and international conference proceedings. His current research interests include multiobjective algorithms for dynamic scheduling and resource management in weather-impacted environments. Mr. Sidoti was an Invited Speaker at NATOs Decision Support and Risk Assessment for Asset Planning 2015 Workshop.

GOPI VINOD AVVARI received the B.Tech. degree in electronics and communication engineering from Acharya Nagarjuna University, Guntur, India, in 2011. He is currently pursuing the Ph.D. degree with the Department of Electrical and Computer Engineering, University of Connecticut, Storrs, CT, USA, under the advisement of Dr. K. Pattipati. His current research interests include developing proactive decision support tools for unmanned aerial vehicles (UAVs); multiobjective algorithms with weather-impacted environments for ship, helicopters, and UAV; stochastic control moving horizon problem involving dynamic programming and combinatorial optimization; OCV-SOC modeling; estimating capacity and resistance of Li-ion batteries; developing chemistry adaptive BFG; and hardware-in-loop tests to validate battery fuel gauge algorithms.

PUJITHA MANNARU received the B.E. degree in electronics and communications engineering from the PES Institute of Technology–Bangalore South Campus, Bengaluru, India, affiliated with Visvesvaraya Technological University, India, in 2013. She is currently pursuing the Ph.D. degree in systems engineering with the University of Connecticut, Storrs, CT, USA. Her research interests are in the areas of applications of signal processing, machine learning, and psychophysiological measures to proactive decision support tools and human–machine systems.

DIEGO FERNANDO MARTINEZ AYALA received the B.S. degree in electrical engineering from the Industrial University of Santander (Universidad Industrial de Santander), Colombia, in 2007, and the M.Sc. degree in electrical engineering from the University of Connecticut, in 2014, where he is currently pursuing the Ph.D. degree in electrical and computer engineering. His current research interests include the development of models and methods for synthesizing adaptive organizations and for planning in dynamic and uncertain environments, multi-objective optimization techniques for engineering applications, development of models of distributed information processing in human teams, Bayesian networks, Markov processes, and pattern recognition in the context of big data analysis, as well as renewable energy technologies with SMART grid applications.

KRISHNA R. PATTIPATI (S'77–M'80–SM'91–F'95) received the B.Tech. degree (Hons.) in electrical engineering from IIT Kharagpur, India, in 1975, and the M.S. and Ph.D. degrees in systems engineering from the University of Connecticut (UConn), Storrs, in 1977 and 1980, respectively. From 1980 to 1986, he was with ALPHATECH, Inc., Burlington, MA, USA. Since 1986, he has been with UConn, where he is currently a Board of Trustees Distinguished Professor and UTC Professor of Systems Engineering with the Department of Electrical and Computer Engineering. He has served as a Consultant to ALPHATECH, Inc., Aptima, Inc., and IBM Research and Development. He is the co-founder of Qualtech Systems, Inc., which is a small business specializing in intelligent diagnostic and prognostic software tools. His current research interests are in the areas of adaptive organizations for dynamic and uncertain environments, data-driven, model-based and knowledge-based diagnosis and prognosis techniques for complex system monitoring, dynamic sensor management for search and rescue, and incremental learning. Dr. Pattipati was selected by the IEEE Systems, Man, and Cybernetics Society as the Outstanding Young Engineer of 1984, and he received the Centennial Key to the Future award. He was elected as a fellow of the IEEE in 1995 for his contributions to discrete-optimization algorithms for large-scale systems and team decision making. He was the co-recipient of the Andrew P. Sage award for the Best SMC Transactions Paper for 1999, the Barry Carlton award for the Best AES Transactions Paper for 2000, the 2002 and 2008 NASA Space Act Awards for A Comprehensive Toolset for Model-based Health Monitoring and Diagnosis, the 2003 AAUP Research Excellence Award, and the 2005 School of Engineering Teaching Excellence Award from UConn. He was also a recipient of the best technical paper awards at the 1985, 1990, 1994, 2002, 2004, 2005, and 2011 IEEE AUTOTEST Conferences and at the 1997 and 2004 Command and Control Conferences. He has served as the Editor-in- Chief of the IEEE TRANSACTIONS ON SYSTEMS, MAN, AND CYBERNETICS: PART B-CYBERNETICS in 1998–2001, Vice President for Technical Activities of the IEEE SMC Society in 1998–1999, and Vice President for Conferences and Meetings of the IEEE SMC Society in 2000–2001.

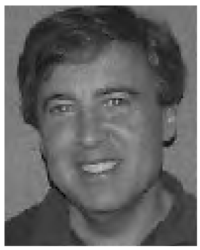

DAVID L. KLEINMAN (S'62–M'68–SM'93–F'94) received the B.E.E. degree from The Cooper Union, New York, NY, USA, in 1962, and the M.S. and Ph.D. degrees in electrical engineering (control systems) from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 1963 and 1967, respectively. He is a part-time Research Professor with the Department of Information Sciences, Naval Postgraduate School, Monterey, CA, USA, and is also a Professor Emeritus with the Department of Electrical and Computer Engineering, University of Connecticut (UConn). From 1967 to 1971, he was with Bolt Beranek and Newman, Cambridge, where he pioneered in the application of modern control and estimation theory to develop and validate an analytical model for describing human control and information processing performance in manned vehicle systems. From 1971 to 1973, he directed the East Coast Office of Systems Control, Inc., where he led applied research projects in both manual and automatic control. From 1973 to 1994, while he was a Professor at UConn, he directed the CYBERLAB, which is a laboratory for empirical research in cybernetic systems. He has authored or co-authored over 200 journal and conference papers, many of which are in the area of military command and control. He was the co-founder of ALPHATECH, Inc.; Qualtech Systems, Inc.; and Aptima, Inc. He continues to serve on the Board of Directors of the latter two companies. His current research is in the areas of team decision making/coordination in dynamic task environments and structural adaptation processes in human organizations. Dr. Kleinman has been the Vice President of Finance and Long-range Planning for the IEEE-SMC Society. He was the Program Chairman of the 1989 IEEE International Conference on Systems, Man, and Cybernetics and on the 1990 JDL Symposium on Command and Control Research. In 1993, he was named an IEEE fellow for his work on human modeling.