Context-Driven Proactive Decision Support for Hybrid Teams

Manisha Mishra, Pujitha Mannaru, David Sidoti, Adam Bienkowski, Lingyi Zhang, Krishna R. Pattipati

■ *A synergy between AI and the Internet of Things (IoT) will significantly improve sense-making, situational awareness, proactivity, and collaboration. However, the key challenge is to identify the underlying context within which humans interact with smart machines. Knowledge of the context facilitates proactive allocation among members of a human–smart machine (agent) collective that balances autonomy with human interaction, without displacing humans from their supervisory role of ensuring that the system goals are achievable. In this article, we address four research questions as a means of advancing toward proactive autonomy: how to represent the interdependencies among the key elements of a hybrid team; how to rapidly identify and characterize critical contextual elements that require adaptation over time; how to allocate system tasks among machines and agents for superior performance; and how to enhance the performance of machine counterparts to provide intelligent and proactive courses of action while considering the cognitive states of human operators. The answers to these four questions help us to illustrate the integration of AI and IoT applied to the maritime domain, where we define context as an evolving multidimensional feature space for heterogeneous search, routing, and resource allocation in uncertain environments via proactive decision support systems.*

Rapidly changing patterns in today's world impose real-time decision-making requirements in many complex organizations, ranging from maritime establishments to agile manufacturing systems and commercial enterprises. One of the key trends in maritime operations is the pervasive use of smart machines (for example, unmanned aerial and underwater vehicles) for countersmuggling, search and rescue operations, and battle management, to name a few. The primary reasons for the use of unmanned vehicles include their operability from remote locations, ultralong endurance, and high-risk mission acceptance. Additionally, these smart machines can be made smaller, agile, and more economical than their manned counterparts. With rapid advances in AI, smart machines and agents are becoming more autonomous (that is, they can both respond to human commands and operate independently) and require only intermittent human intervention to keep them aligned with human intentions. In March 2018 (Wakabayashi 2018), a pedestrian was hit and killed by a self-driving car (with an emergency driver behind the wheel), indicating the need for a human driver to intervene in a timely manner to avoid mishaps. It was later found that the emergency braking system in the car was disabled and the driver was streaming a TV show before the incident. This mishap could have been avoided had the emergency brakes been enabled and had the driver been paying attention. This incident illustrates the need to make machines more intelligent so that they can mimic the human thought processes of understanding the environment and can provide appropriate alerts for interfacing with the human when attention is warranted or required for not only unexpected but uncertain situations as well.

In fast-paced decision-making environments, it is well established that human performance becomes suboptimal when the workload is too high, as well as when it is too low. Channelized attention has been implicated in numerous operational mishaps. According to a 2012 US unmanned aerial vehicle (UAV) report to Congress, 68% of UAV accidents are attributable to human error. To prevent such mishaps, there is a need to develop hybrid teams, comprising humans and smart machines, agents, and devices, for proactive decision making in complex and uncertain environments. A human–smart machine collective has the potential to exploit the strengths of rapid computation, communication and control capabilities of smart machines and agents, and the generalization, learning, and adaptation proficiencies of humans for improved sensemaking, situational awareness, collaboration, and coordination when compared with teams composed of only humans or machines.

With the advent of the Internet of Things (IoT), it is possible to have humans and smart machines connected anytime, anywhere, with anything and anyone, ideally using any path or network and any service. IoT allows seamless interoperability and connectivity among smart machines, facilitating rapid responses to user requests and potentially ubiquitous access to information.

One of the key technical challenges in integrating IoT, decision support systems (DSSs), and human teams is identifying the context where a set of smart machines needs to be connected so that the right information from the right source and right context are conveyed to the right user at the right time for the right purpose (Smirnov 2006). Figure 1 depicts the data-to-decisions process, including the human-machine interaction. IoT plays a crucial role in real-time decision making as it facilitates communication among the stages of the process: (1) data sensing, (2) data analytics to determine the current context, (3) diagnosis of root causes of contextual changes and prediction of the future evolution of context for relevant courses of action, (4) dissemination of contextual information to the correct user, and (5) decision making and action validation. Figure 1 also depicts how AI in smart machines evolves from assisted intelligence to augmented intelligence to autonomous intelligence. Assisted intelligence is provided for repetitive and routine tasks, where the nature of tasks does not change. In the case of augmented intelligence, tasks are dynamic, and humans and machines learn from each other. Finally, when the machines are capable of understanding the situation and taking relevant actions, they require minimal human intervention to achieve the human's intent. Here, we focus on augmented and autonomous intelligence in the domain of maritime decision making. In this article, smart machines are DSSs that interact with human decision makers to develop proactive courses of action and facilitate what-if contextual analysis for robust and adaptive decision making.

Literature Review

The idea of machines performing tasks that typically require humanlike understanding has been around for more than 60 years. With cloud computing and the deluge of data, machine learning techniques, especially deep-learning methods, have evolved rapidly since 2012 to demonstrate human-level capability; for instance, machines are able to recognize objects in digital images, even when the objects are represented in unusual contexts, such as varying backgrounds and levels of light. Articulate AI, which combines deep learning, unsupervised and semisupervised learning and causal reasoning, and the generalization that is innate to humans (that is to say, domain knowledge, how the constituent parts fit together, causal relationships), is a critical need for mixed-initiative proactive decision making. Creativity — articulating and finding the problem, while seeing problems that are not there through look-ahead reasoning, anticipatory behavior, and asking what-if questions — will always be the human frontier. The human planner is indispensable, working in an interactive

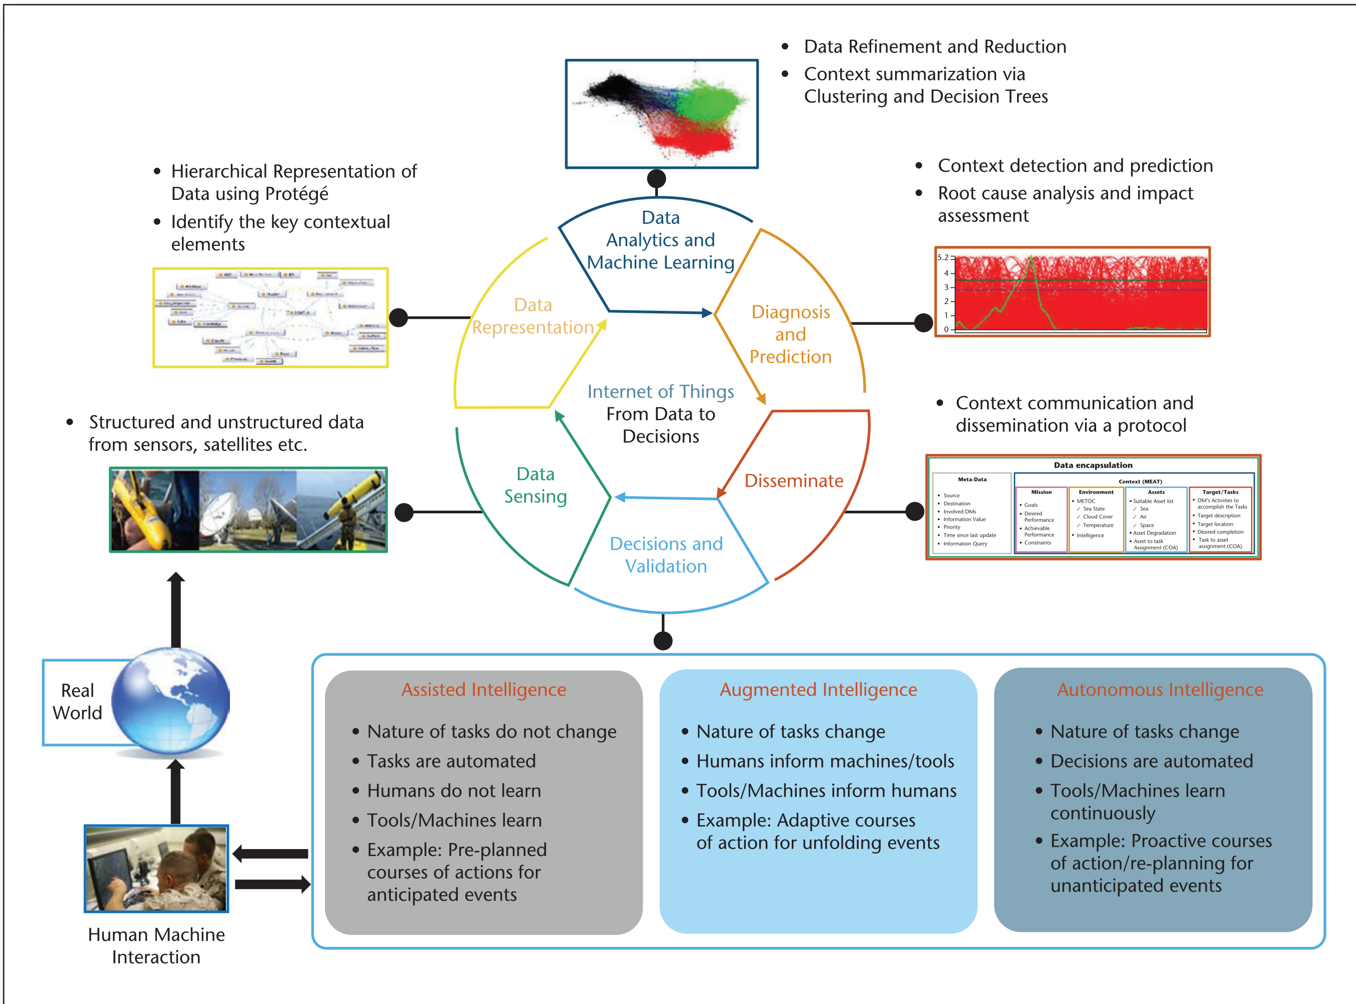

Figure 1. Proactive Decision Making.

Central to the Internet of Things, proactive decision making involves (1) sensing data from heterogeneous sources via sensors, (2) representing data and their interdependencies via Protégé (Mishra et al. 2017b) for determining context, (3) processing context via analytics and machine learning, (4) context detection and prediction via root cause analysis, (5) context dissemination via communication protocols, and (6) decision and validation via real-time applications.

way with AI, because unexpected machine system states occur. Table 1 provides the relative strengths and limitations of humans and machines, sometimes referred to as what humans are better at and what machines are better at (HABA-MABA), as defined by Fitts in the 1950s (Fitts et al. 1951). Since then, a series of technical studies have been conducted on automation and on the formation of collaborative teams of humans and machines (Chapanis 1965; Muir 1987; Rasmussen 1983).

Table 2 lists the levels of automation of decision and action selection in a human-machine collective. Parasuraman. Sheridan, and Wickens (2000) provided a four-stage model of human-automation interaction, comprising information acquisition, information analysis, decision selection, and action implementation guidance on the type and level of automation in a system. We believe that today, human decision makers are an integral part of decision making in human-machine systems, and thus our applications deal mostly with levels 2 through 5.

Rasmussen (1983) categorized human behaviors into (1) skill-based sensory-motor actions that are highly automatic and acquired after some period of training; (2) rule-based actions guided by subroutines, stored rules, and if-then-else structures; and (3) knowledge-based actions, where mental models built over time aid in the formulation and selection of plans for an explicit goal. Cummings (2014) added a fourth behavior, that of expertise, the highest level of cognitive control achieved from significant experience in a particular field in the presence of uncertainty. Table 3 describes the varying degrees

Attribute	Machine	Human
Speed	Superior	Comparatively slow
Power	Output superior in level and consistency	Comparatively weak
Consistency	Ideal for consistent, repetitive action	Unreliable learning and fatigue are factors
Information capacity	Multichannel	Primarily single channel
Memory	Ideal for literal reproduction, restricted, and formal access	Better for principles and strategies, versatile and innovative access
Reasoning computation	Deductive, tedious to program, fast and accurate, and poor error correction	Inductive, easier to program, slow, accurate, and good error correction
Sensing	Good at quantitative assessment, poor at pattern recognition, poor judgment ability	Wide ranges, multifunction, good judgment capability
Perceiving	Copes with variation poorly, susceptible to noise	Copes with variation better, susceptible to noise

Table 1. Comparisons of Humans and Machines.

High
9. Informs the human only if it, the machine, decides to
8. Informs the human only if asked
7. Executes automatically, then necessarily informs the human
6. Allows the human a restricted time to veto before automatic execution
5. Executes a suggestion if the human approves
4. Suggests one alternative
3. Narrows the selection down to a few alternatives
2. Offers a complete set of decision/action alternatives
1. Offers no assistance: the human must make all decisions and actions.
Low

Table 2. Levels of Automation of Decision and Action Selection.

of automation for tasks of various information processing behaviors. Figure 2 provides a few examples of tasks and the corresponding degree of automation required (Cummings 2014).

We briefly review three concepts, namely, context, decision support, and IoT, related to this article.

Context

The definition and representation of context enables related information to be attached to data for later retrieval and decision making. The word *context* has Latin roots: *con* means to join together or to weave together, and *texere* means to weave or to make, implying weaving together the circumstances that form the setting of an event or scenario. In computing, for example, context has been referred to as the location and identities of nearby people and objects, and changes to those objects; knowledge about the user's and the machine's state, including surroundings, situation, and location; and emotional state of the user, focus of attention, and information the user is attending to. Further, context has been decomposed into various categories, such as computing context, user context, physical context, time context, and cognitive context. Although there is no unified definition of context, one of the most celebrated works in human-computer interaction, by Abowd et al. (2001), formally defined context as

> any information that can be used to characterize the situation of an entity (for example, a person, place, or object) considered relevant to the interaction between a user and application, including the user and applications themselves.

Context awareness — the ability of a system to provide relevant information or services to users via the utilization of contextual information, where relevance depends on the user's task — has become a key factor for successful integration of human-machine systems. Further, context awareness has been categorized as active and passive: active context awareness enables an application to automatically adapt to a discovered context by changing the application's behavior, whereas passive context awareness enables the application to present the new or updated context to an interested user or make the context persistent for the user to retrieve later. We define context as a multidimensional feature space, which dynamically evolves with time.

DSSs

A DSS is an information processing system embedded into an organizational decision-making system to aid the decision maker in solving unprogrammed, unstructured, or semistructured problems (Bonczek, Holsapple, and Whinston 2014). Decision making was defined by Fitts et al. (1951) in an air traffic control design scenario as the basic human action involved in the human-machine system — the choice by a pilot or controller of a particular course of action (COA) from among the alternatives available at a particular moment. Decision making was further explored from a system's perspective by Horvitz (1999); in that work, the system makes decisions as a function of an inferred probability of performing specific operations. DSSs are categorized as model driven, communications driven, data driven, document driven, knowledge driven (event driven), or context driven. Model-driven DSSs emphasize access to and manipulation of a quantitative model, making it the dominant component in the DSS architecture. Communications-driven DSSs derive their functionality from communications and information technologies within the system to support shared decision making. The functionality of data-driven DSSs results from access to and manipulation of a large database of structured data. Document-driven DSSs integrate a variety of storage and processing technologies to provide sophisticated document retrieval and analysis to support decision makers. Knowledge-driven DSSs suggest or recommend actions based on knowledge that has been stored using AI or statistical tools, such as case-based reasoning, rules, frames, and Bayesian networks. Finally, context-driven DSSs derive their functionality from dynamically integrating knowledge, that is, (1) relevant to the mission, environment, assets, threats, or tasks including the decision maker's activities; (2) informed by up-to-date data sources; and (3) congruent with the work flow and the individual decision maker's role in mission, workload, and expertise.

IoT

The IoT integrates the physical world with the virtual world of the Internet. Objects such as car, house, clothes, and refrigerator are electronically tagged with important information and connected to the Internet through remote, contactless technology. These things are equipped with sensors that collect, optimize, and send real-time data to a decision maker, other devices in the network, or hybrid teams in the human-machine system. Context awareness plays an integral role in deciding on the information to be collected, services to be executed automatically, and sensor data to be fused within the network (Perera et al. 2014). For example, a shipping company can offer context-aware functionalities via the fusion of relevant data, such as ship location, speed, and cargo temperature, to optimize maintenance costs and route planning (Tracy 2016).

Cognitive Behavior/Task	Degree of Automation
Skill-based	Best candidate for automation, assuming reliable sensors for state and error feedback
Rule-based	Possible candidate for automation, if rule set is well established and tested
Knowledge-based	Some automation can be used to help organize, filter, and synthesize data
Expertise-based	Human reasoning is superior, aided by automation as a teammate

Table 3. Degree of Automation as a Function of Desired Behavior.

Technical Challenges

In this section, we list challenges in proactive decision making by hybrid teams, within the maritime domain, focusing on augmented and autonomous intelligence. The technical challenges can be categorized depending on the stage of data to decision-making process. They include context representation challenges, context determination and data analytics challenges, context communication challenges, and decision and validation challenges.

Context representation challenges include (1) how to define, represent, identify, and characterize evolving contextual elements, including human-machine-task interactions, and (2) how to summarize contextual attributes (for example, speed, capability, sweep (search) width or detectability index of assets, environmental parameters, cognitive state of humans) without making a problem computationally expensive.

Context determination and data analytics challenges include (1) how to determine context and contextual changes within an uncertain environment and (2) how to predict the evolution of context and develop flexible (for anticipated and evolving events) as well as agile (for unanticipated events)

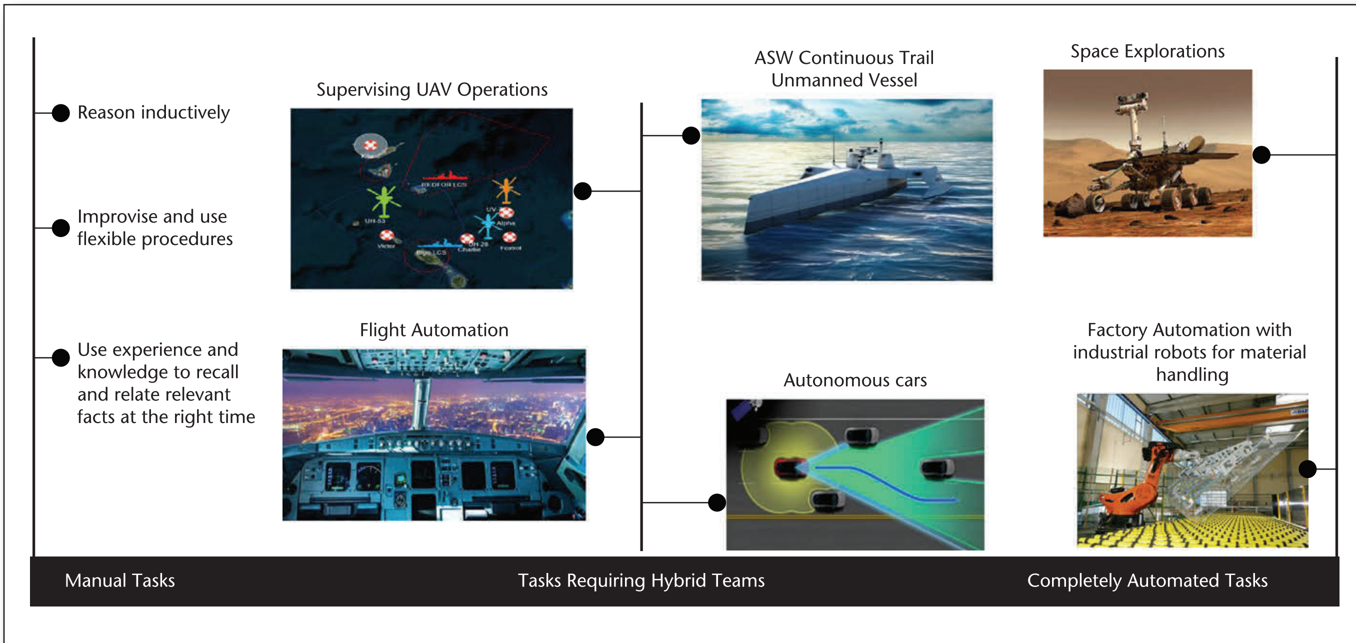

Figure 2. Manual Tasks, Tasks Requiring Hybrid Teams, and Completely Automated Tasks.

COAs to achieve resilience in dynamic and uncertain environments.

Context communication challenges include (1) how to communicate context among hybrid team members in a timely manner and (2) how to minimize communication costs among different team members.

Decision and validation challenges include how to explore methods for instantiating alternative COAs consistent with the needs of system goals when context changes are detected and root causes inferred and how frequently replanning should occur (for example, as events emerge or at a predefined frequency). A third challenge is how best to formulate and allocate tasks among the available assets and machines and when to allow human intervention for superior system performance — in particular, how automation makes its behavior known to human agents, and how easily and efficiently it is directed. Part of this challenge is determining when and at what level the automation should be enabled (adaptive and adaptable automation differ according to whether the machine or human, respectively, possesses responsibility for controlling the level of automation). Finally, how should a scalable human-machine collective be integrated for a range of missions? In particular, how can the human-to-machine ratio be reduced to reduce the cost of manpower and training, and how can human workload be accurately estimated and balanced while maintaining overall human-machine capabilities in dynamic and uncertain mission environments?

Addressing these challenges is beyond our scope; however, we focus on (1) representation of the key elements of a hybrid team, (2) context determination and analysis, (3) dynamic task allocation and methods to reduce human workload in time-critical mission environments, and (4) decision support tools (that is, machines) to represent, identify, predict, and learn new context and provide appropriate COAs.

Modeling Hybrid Teams

The triad of human-machine-task interactions is best represented using graphical models learned from data. The models should include task variables (task types, amount of work, task disturbances, dependencies among tasks), team variables (capabilities/expertise in terms of work rates, communication/coordination, learning), and quality of interaction (internal-external work, task time, task accuracy). Data-enabled predictive models can provide insights into high-performing human-machine collectives. For example, if the mission is a notionalized point in multidimensional task work space, with each human-machine capability as a vector, the work vector and the weighted combination of capability vectors, with weights being effort probabilities over time, must be congruent. The rationale for alignment stems from queuing network theory and theory of constraints: match service rates at agent nodes (human or machines) to get maximum throughput.

The multidimensional operational context in the maritime domain comprises an external context in terms of mission, environment, assets, and threats or tasks, and an internal context composed of the cognitive states of human and decision maker (see figure 3). Each MEAT-H (mission, environment, assets, threats/tasks–humans) element may have several subelements with associated states where level of specificity may vary with decision-maker roles. Examples of missions are antisubmarine warfare, countersmuggling operations, and UAV coordination in heterogeneous environments (Mishra et al., forthcoming; Mishra et al. 2014; Sidoti et al., forthcoming). Each mission is characterized by goals, desired performance, achievable performance, and constraints.

Examples of environment elements are cloud cover, sea state, precipitation, salinity, and temperature, each with possible multiple states.

Examples of assets are sea, air, space, and land assets, with subelements such as frigates, high- or medium-endurance cutters, and P-3s, where subelement states may include an asset's availability (available, unavailable) or crew endurance (in hours).

Finally, note that the T in the MEAT-H framework may refer not only to the threats or tasks to be done but also to the decision maker's activities to accomplish them (namely, the work flow). Some examples of threats or tasks are the interdiction of drug smugglers, protection of high-priority maritime vessels, and reconnaissance, depending on the operational mission context. In this article, a mission may consist of subgoals, referred to as threats or tasks. A proactive DSS considers the decision makers' workloads, time pressures, and roles in the determination and communication of relevant information for effective mission performance. Once context is represented, it must be analyzed via approaches like event trees or decision trees as discussed in supplementary material by Mishra et al. (2018).

Adaptive model-based approaches to context-based mission representation include graphical models like dependency graphs (digraphs), Petri nets, multifunctional flow graphs, action-goal attainment graphs, hidden Markov models (HMMs), coupled HMMs, factorial HMMs, and dynamic hierarchical Bayesian networks (DHBNs). These graphical models are consistent with planning as probabilistic inference, advocated by cognitive scientists and neuroscientists. In their view, a decision-making agent has an internal generative model of the future as a joint probability distribution over actions, outcome states, and rewards and costs. This probabilistic generative model facilitates computationally efficient algorithms for perception and action selection. For example, each node in a DHBN model represents contextual elements (user intent, environment, asset, threat or task, human cognitive context), and edges denote the dependencies among pairs of contextual elements. DHBN-based contexts represent these characteristics: (1) finite, but large number of states; (2) multistage representations, with decisions made at each stage's beginning; (3) stochastic effects generated from action execution; (4) function execution in a particular state generates one of possible numerous states with associated probabilities reflecting unforeseen external events (for example, unanticipated threats, weather, terrain); (5) complete or partial observation of true states of environments at any stage; and (6) goal-directed decision making.

Context Determination, Root Cause Analysis, and Inference

Any decision originates in a dissatisfaction from differences between the current state and a more desirable state not yet existing (Pomerol 1997). Causal models use differences between model predictions of context and operational data to detect context changes, triggering (1) root cause analysis (diagnosis) of context changes (that is, which contextual elements caused change), (2) projection of the impact of changed context on mission goals (what could happen with associated uncertainty), and (3) replanning strategies to proactively explore decision alternatives that exploit potential opportunities or mitigate negative consequences of changed contexts (what COAs need to be taken).

Statistical hypothesis testing and machine learning techniques detecting context changes and neurodynamic programming-based algorithms suggest COAs commensurate with changed context. Update of the DHBN model structure via planner-in-the-loop active learning enables the discovery, labeling, and incorporation of new context with associated feature data. Quantification of value of information, coupled with free energy–based inference and decision making, enables exploration-exploitation trade-offs inherent in sequential decision making under uncertainty. Recently, a theory has been proposed to suggest that agents, for example, biologic systems, such as a cell or brain, adapt to their environments by reducing the information-theoretical quantity known as variational free energy. This theory, called the free energy principle (Friston 2010) or maximum entropy reinforcement learning (Levine 2018), brings information-theoretic, Bayesian, neuroscientific, and machine learning approaches into one framework by formalizing that decision makers (more generally, agents) reduce free energy in three ways:

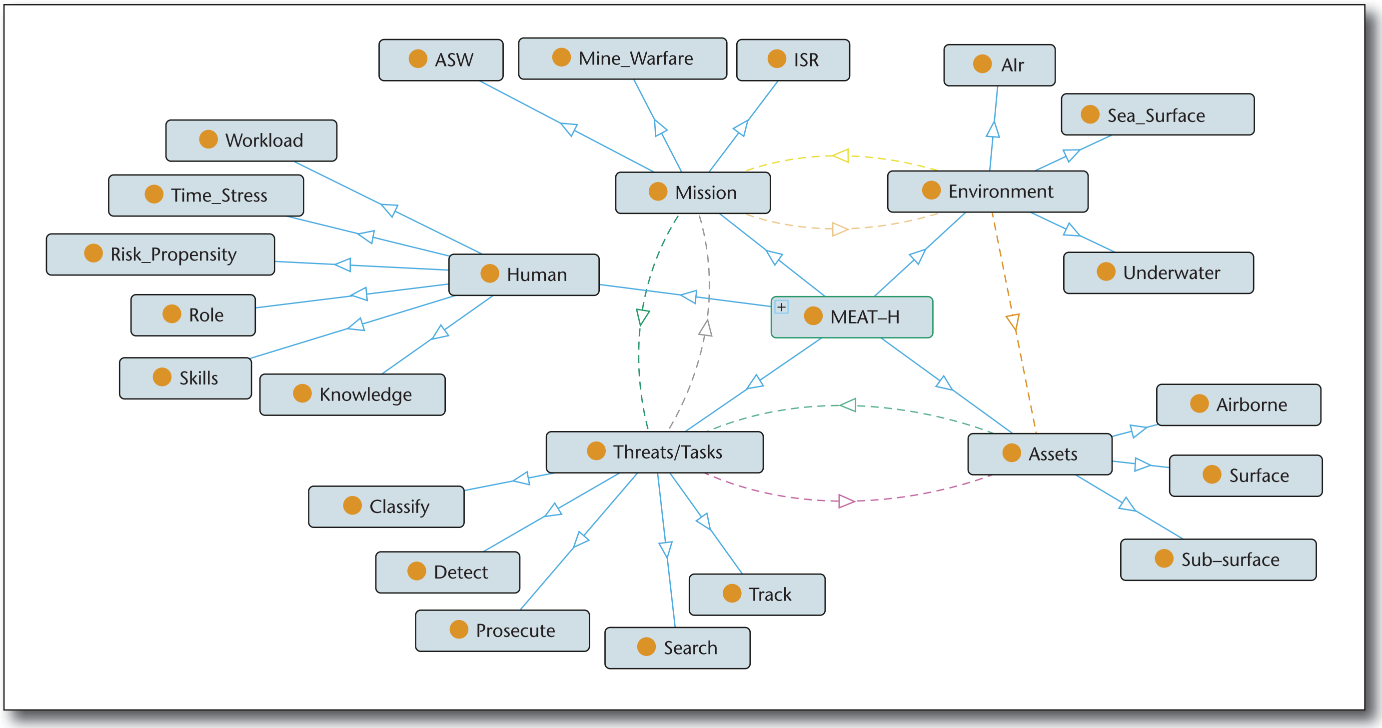

Figure 3. Representing Contextual Elements.

Contextual elements (MEAT-H) are represented in hybrid teams for maritime decision making via Protégé (Mishra et al. 2017b). Solid arrows represent the hierarchy, whereas dashed arrows denote relationships between elements. In this figure, mission elements involve antisubmarine warfare, mine warfare or intelligence, surveillance, and reconnaissance operations.

(1) by changing sensory input (information seeking action selection), (2) by changing predictions of sensor inputs (perceptions, beliefs), and (3) by changing the model of the decision maker's team structure and coordination (learning). Variational free energy is a function of sensory outcomes (data) and probability density over their (hidden) causes (true world states or context). This function is an upper bound on surprise, a negative log of model evidence representing the difference between agent predictions about sensory inputs, and observations or data encountered. Indeed, differences between variational free energy and surprise is Kullback–Leibler divergence between agent beliefs about context (called the recognition density) and the joint density of context and data given the agent model (the generative density). Because the long-term average of surprise is entropy, an agent acting to minimize free energy will implicitly place an upper bound on the entropy of outcomes or sensory states sampled. Consequently, the free energy principle provides a mathematical foundation to explain how agents maintain order by restricting themselves to a limited number of perceived high-probability and high-utility (context, action) pairs. This restriction gives a formal mechanism to inference and decision making, where multiple agents operate autonomously, coordinate among themselves, and resist disorder (Levchuk et al. 2018).

Context inference compares model predictions of outcomes with expected (desired) outcomes; these deviations form the basis for nonnormalcy detection, active learning, and subsequently predictions based on updated context model. Because context inference is a maximum a posteriori estimation problem on a DHBN, we convert the DHBN model into a factor graph of data variables, context factors, and decision makers or agents. Algorithms for dynamic context inference using factor graphs of data variables and context factors (assuming each factor is associated with an agent) include generalized belief propagation algorithms, decomposition algorithms, and clustering approaches.

Generalized belief propagation algorithms (Yedidia, Freeman, and Weiss 2005) include a combination of coordinate descent, Lagrangian relaxation and Viterbi decoding algorithms developed for coupled HMMs (Zhang et al. 2013), and semisupervised (active learning–based) clustering algorithms. Because exact computation of posterior context distribution is NP-hard, an approximate solution is produced by the max-product belief propagation algorithm, which is the Bethe approximation of free energy function. Max-product belief propagation computes max-marginal distributions by iteratively passing belief messages between variable (data) and factor nodes in factor graphs.

Decomposition algorithms decompose the inference problem into decoupled subproblems, one for each contextual element (pattern of mission, environment, asset, threat, or task); the subproblems, solved in parallel, are coordinated by updating Lagrange multipliers and are iterated until convergence. Each subproblem corresponds to finding optimal context-state sequences, solved using the Viterbi decoding algorithm (George et al. 2017). This approach is suited for distributed, asynchronous implementations.

Clustering each contextual element and assigning new data to known contexts or classifying it as unknown. These unknown contexts are labeled by decision makers with an updated context model used for predicting the impact of context changes on mission performance via Q functions that approximate the cost-to-go for (context, action) pairs.

Human-Machine Task Allocation

Task allocation assigns mission tasks to humans or machines in real time by monitoring agent activities and identifying tasks best performed by humans or machines based on triggering contexts. Tasks assigned to machines or agents require multiobjective optimization algorithms to meet desired goals. Multiobjective optimization seeks a Pareto front, representing all trade-offs among mission planning objectives. A human decision maker can explore and decide which available trade-off works best via an adaptive scatter or gather or a baseline web search interface for exploratory or lookup-type query tasks, respectively. Multiobjective planning (for example, resource allocation, routing, scheduling) can be viewed as a moving horizon stochastic control problem. The key is to optimize or equilibrate mission objectives, subject to realistic constraints such as weather, asset capabilities (for example, range, speed), and asset assignment (for example, coordination among multiple sensors for improving situational awareness). These problems are computationally intractable because of multiple conflicting objectives and uncertain mission environments. Key concepts include problem decomposition, approximate dynamic programming, limited search, domain-specific constraints, and scalable Pareto optimization approaches that exploit efficient data structures and multiobjective A* algorithms for dozens of objectives or more.

Additionally, uncertainty in maritime missions can range from situations where probability distributions over outcomes is known, partially known (for example, they belong to a family of distributions over outcomes), and unknown, requiring online learning. In such cases, scenario-based uncertainty management approaches can be applied, including robust, flexible, and agile methods.

Robust decision making seeks to manage uncertainty via minimizing variability in the expected risk or reward, minimizing maximum risk, or maximizing minimum reward. Methods based on *m*-best solutions coupled with robust design techniques, cone programming, and efficient propagation of Pareto solutions in search spaces are representative of these methods.

Flexible decision-making methods adapt to uncertain contexts via enumerating or brainstorming potential event sequences a priori, conducting what-if analyses, and preplanning response policies. These methods adapt to expected scenarios by recognizing critical events signaling context change. Typical methods in the context of planning include FRAG plans (fragment-based branches and sequels) and conformant plans (Kurien, Nayak, and Smith 2002). Decision-directed open-loop optimal feedback, Q-learning, and rollout strategies belong to this category.

Agile decision-making methods adapt to uncertain and unexpected contexts by learning (online) an updated model of decision environments while hedging against uncertainty by trading off exploration versus exploitation (also known as dual control or probing and caution in the stochastic control literature [Powell 2007; Sutton and Barto 1998]). Typical methods in the context of planning include coupling active learning algorithms with moving horizon planning, certainty equivalence, open-loop optimal feedback, free energy optimization, and other approximate dynamic programming techniques.

Graphical models, coupled with concepts from multiobjective optimization, provide an elegant framework for adaptive distributed task allocation in multiagent systems, like human-machine collectives. A basic formulation of human-machine task allocation is provided in supplementary material (Mishra et al. 2018). This model characterizes skills of humans and machines in terms of quality (task accuracy) and cost (measured in terms of time, processing and coordination workloads, and other economic factors) and recognizes the many different ways that humans and machines may cooperate.

Application of Context-Driven Decision Making in Hybrid Teams

Maritime operations often require integrated, interconnected platforms for rapid decision making and planning in uncertain and dynamic environments. The US navy envisions a networked battlespace linking target information from heterogeneous sources, such as aircraft, ships, and underwater vehicles, to keep pace with potential adversaries. A networked battlespace facilitates the flow of relevant information among heterogeneous sensors at geographically disparate locations via cloud-based IoT, allowing heterogeneous assets (UAVs, ships, and submarines) to access a range of targeting information for carrying out operations like integrated air and missile defense, countersmuggling operations, and multiobjective ship routing. The exchange of relevant context or information among sensors and platforms allows for timely fusion by analysts to develop proactive COAs using decision support tools. To demonstrate the relevance of IoT and AI in a networked battlespace environment, we discuss two areas of operations and the concomitant decision support tools for providing augmented and autonomous intelligent courses of action for (1) multidomain battlespace management (MDBM) and (2) supervisory control operations user testbed (SCOUT). The latter application can be found in supplementary material by Mannaru et al. (2017) and Mishra et al. (2018).

MDBM requires situation awareness and advanced battlespace management to place assets where and when needed and to make the (anticipate) → find → fix → track → target → engage → assess cycle faster and more resource efficient (see figure 4). MDBM is a moving-horizon stochastic planning and control loop, where surveillance sensors seek to detect emerging threats or targets (find). Target detection results in an alert (or context change) that initiates a decision-making process to determine whether to prosecute it or to continue surveillance. This decision-making process commences when positive identification of the target is requested and accomplished by the information processing block (fix). The movement of the target is monitored and its track maintained while its desired end effects are confirmed (track). In the targeting phase, available assets are searched for suitable weapon platforms to engage the target based on desired end effects. A collateral damage estimate is performed, and the mission package is reviewed against rules of engagement and submitted for approval. The targeting phase is often the longest because of the large number of constraints to be satisfied. The engage phase commences when ordered by the commander. A brief is drafted and transmitted to platforms that engage the target after acknowledging receipt and comprehension of contents. The loop ends with a battle damage assessment report (assess) and the process repeats. The MDBM problem assigns relevant resources to mission tasks to achieve one or more goals, while satisfying a set of domain constraints, like geographic limitations, meteorological and oceanographic (METOC) priorities, and sequences of mission tasking to address threats.

The solution to the MDBM problem requires networked architectures for easier information flow and decision making. Figure 5 depicts the overall MDBM architecture compatible with the DF-NTC, which includes Docker,[1] a modular, environment-independent way of handling segmented code based on mission context. The Docker software enables transitions of software, algorithms, or database communication networks between systems (or collaborators), while guaranteeing that software will run as expected. The algorithmic decision-making components include the tool for multiobjective planning and asset routing (TMPLAR), conflict identification (CONFIDENT), the courses of action simulation tool (COAST), and asset package selection and planning (APSP). The mission planning elements are as follows: (1) mission includes all MDBM services, for example, TMPLAR, CONFIDENT, COAST, APSP; (2) environment includes bathymetry data and METOC; (3) assets include ship or submarine limits; (4) tasks or threats are input via JavaScript object notation or text files; and (5) human is the decision maker who specifies mission objectives, that is, whether the commander's intent requires conflict identification, traversal between operational areas, search and prosecution of targets, maritime operations center planning, or a combination.

TMPLAR

Navigation under uncertainty involves several contextual elements, like different METOC conditions (ensemble forecasts with varying spatiotemporal uncertainty), evolving multiple objectives (for example, fuel, time, pop-up threats, navigation hazards, bathymetry, depth of fire, avoiding red sensor ranges, extending ship's life, training requirements), and asset condition. A canonical navigation under the uncertainty problem is as follows: Given a graph (for example, grid map probability surfaces), a departure point, and a destination point (including way points), find the set of shortest paths with Pareto efficient costs, where cost may be multiobjective. This problem is the multiobjective shortest path problem under uncertainty with time windows, power plant configurations, speed, and bearing as additional control variables, with time-varying stochastic, nonconvex, and multiattribute costs at nodes and along arcs in the network.

Motivated by practical and economic needs of naval and commercial shipping, we developed TMPLAR to help human planners create ship routes (Sidoti et al. 2016). TMPLAR's goal is to suggest multiple routes

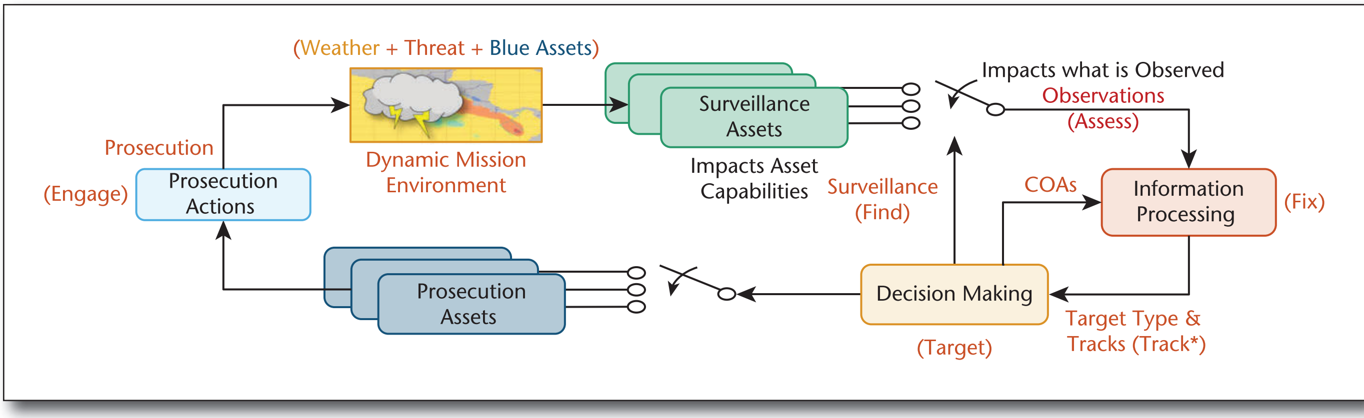

Figure 4. The MDBM Stochastic Control Loop Problem.

(From Mishra et al. 2017; Sidoti et al., forthcoming).

combined with human situational awareness about forecasts (air temperature, sea height, swell height, period, direction, current speed or direction), geographic hazards (light houses, oil rigs, bathymetry, pipelines, undersea cables), asset conditions (hull or propeller fouling, ship dynamics), and threats (torpedoes, sensor ranges) and other uncertainties for a human-machine consensus on routes for one or more ships. A key feature of TMPLAR is the ability to wait at a node, adding a degree of freedom to avoid bad weather or accomplish training. TMPLAR imposes time windows on nodes, that is, the earliest a node can be reached and the latest the ship can depart from a node and still reach its destination given time constraints. Instead of assuming a relative importance of conflicting objectives, TMPLAR finds the set of Pareto efficient solutions, allowing humans to determine which solution is best given trade-offs between objectives. This approach, combined with strategies for finding solutions robust to uncertainty, allows TMPLAR to quickly recommend high-quality routes that humans can select to route a ship. Application of TMPLAR to the El Faro incident (Alvarez, Pérez-Peña, and Robles 2015) illustrated that it potentially could have prevented the loss off the coast of the Bahamas, due to crossing paths with Category 4 Hurricane Joaquin, of a cargo ship heading from Florida to San Juan, Puerto Rico (Bienkowski et al. 2018; Sidoti 2016).

Conflict Identification

Water space planning focuses on route and region deconfliction of submarines and maritime assets to ensure their safe operation (figure 6). Previously, submarine commanders spent hours assigning and deconflicting submarine water space manually. Deconfliction requires intense human supervision and is, therefore, error prone and cumbersome. Our approach to four-dimensional (space and time) trajectory-based conflict detection, where objects are represented as convex regions, like polytopes, ellipsoids, or nonconvex regions, comprises preprocessing to decompose any nonconvex region as the union of several convex regions and a two-phase process for detecting overlaps among convex regions: a broad (coarse) phase and a narrow (fine) phase. For the broad phase, we use R-trees and time parameterized R-trees that scale for tens of thousands of objects. For the narrow phase, we use interior point–based linear programming for polytopes and quadratic programming for ellipsoids. McMenemy (forthcoming) developed a novel pairwise comparison algorithm that determines whether a pair of ellipsoids is overlapping, touching, or separated by exploiting two new Mahalanobis distance-based criteria. These criteria transform the intersection query problem into a least squares minimization over a sphere (a quadratic programming problem). The resulting algorithm is computationally more efficient than previous methods used in video gaming software.

Proactive COAST

Countersmuggling missions involve surveillance operations (to search, detect, track, and identify potential threats) and interdiction operations (to intercept, investigate, and apprehend suspects). Given the probability of activity (POA) surfaces (Hansen et al. 2012) (see figure 7a), which integrate METOC and intelligence information (INTEL) to predict where 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, environment, asset, and threat contexts. Because the mission's geographical environment is large, decision makers assign surveillance assets to specific search regions, and the observations from these assets are processed to characterize target types and trajectories and to correlate contacts

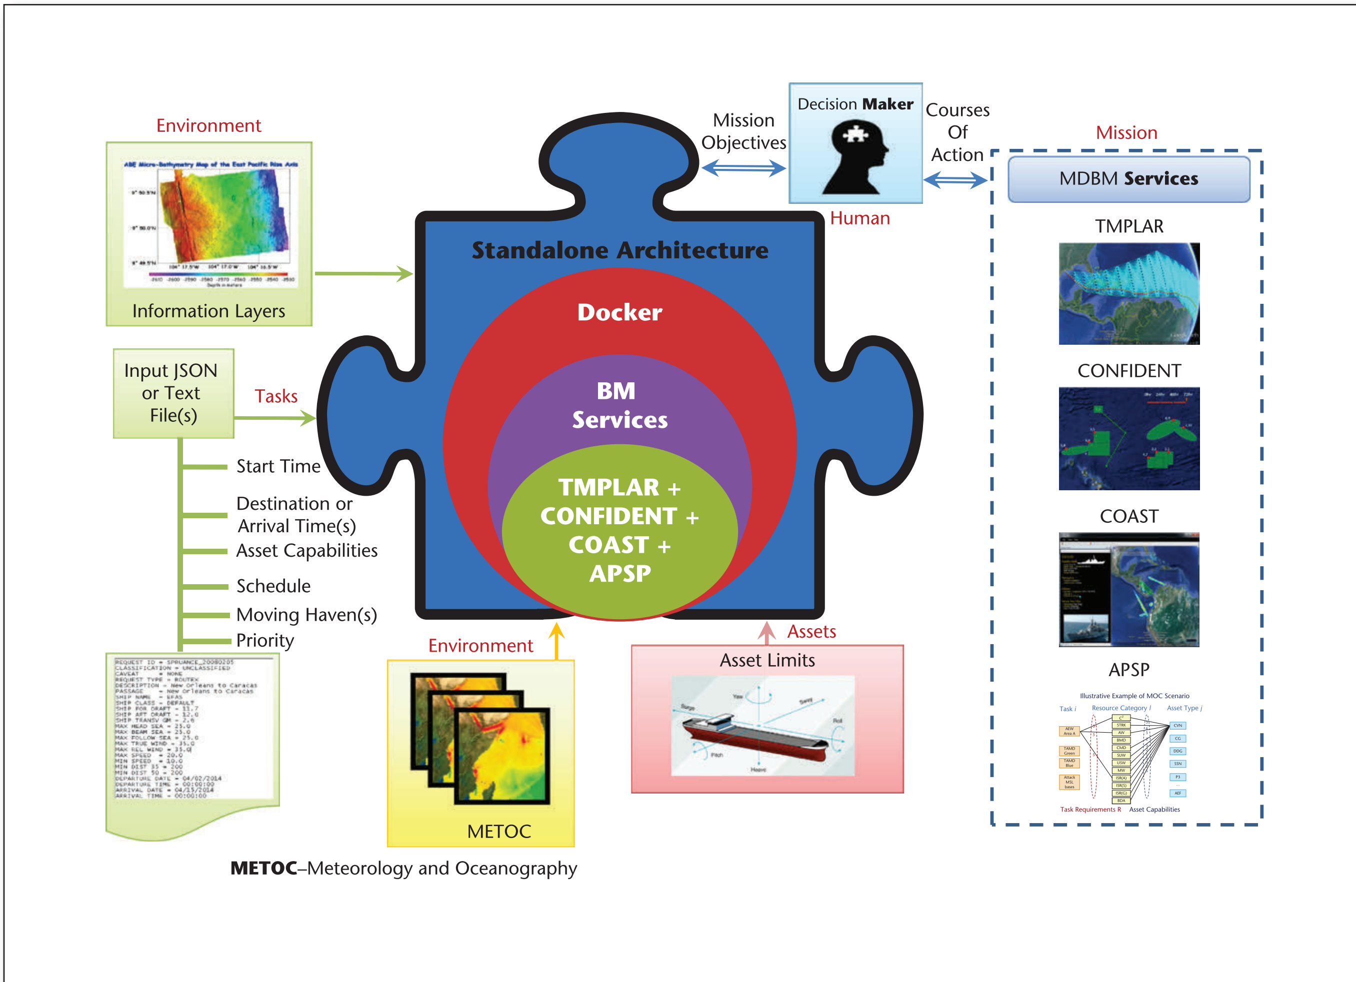

Figure 5. MDBM Architecture.

MDBM with various services: TMPLAR, CONFIDENT, COAST, and APSP.

of interest already located with current INTEL. The newly collected information (for example, INTEL, detections, interdictions, weather data) serve as stimuli (a nonnormal situation when values are out of bounds) for context identification. This information is relayed to the reachback cell in the form of situational reports, using the context protocol (in JavaScript object notation format). The situational reports are then extracted, processed, and aggregated to predict new POA maps for the next planning interval. The predicted POA surfaces are uncertain and can be prioritized based on the weight of contraband to be interdicted or the belief in target INTEL. The context-relevant information gathered by surveillance assets is communicated to interdiction assets using the context protocol for adapting COAs to new contexts. Our approach uses open-loop optimal feedback and consists of asymmetric assignments via a branch and cut algorithm for the surveillance problem and approximate dynamic programming coupled with rollout and Gauss–Seidel techniques for the interdiction problem. These algorithms are embedded in the COAST, an optimization-based decision support tool in a widget format, integrated with Google Earth (see figure 7a).

COAST allows flexible targeting by allocating surveillance and interdiction assets based on mission context. Figure 7b illustrates the prior information regarding target corridors (white rectangles) and related uncertainty (red boxes), analogous to possible locations with possible smuggler activity while taking into account weather and INTEL for the approximate time window of departing a specified port, an input into COAST. In calm weather contexts, the sweep width[2] of surveillance assets is 20 nautical miles (nm). As weather worsens (that is, environmental context changes), sweep width reduces to 2 nm (change in asset performance models), affecting target detection probability of surveillance assets (shown in figures 7c

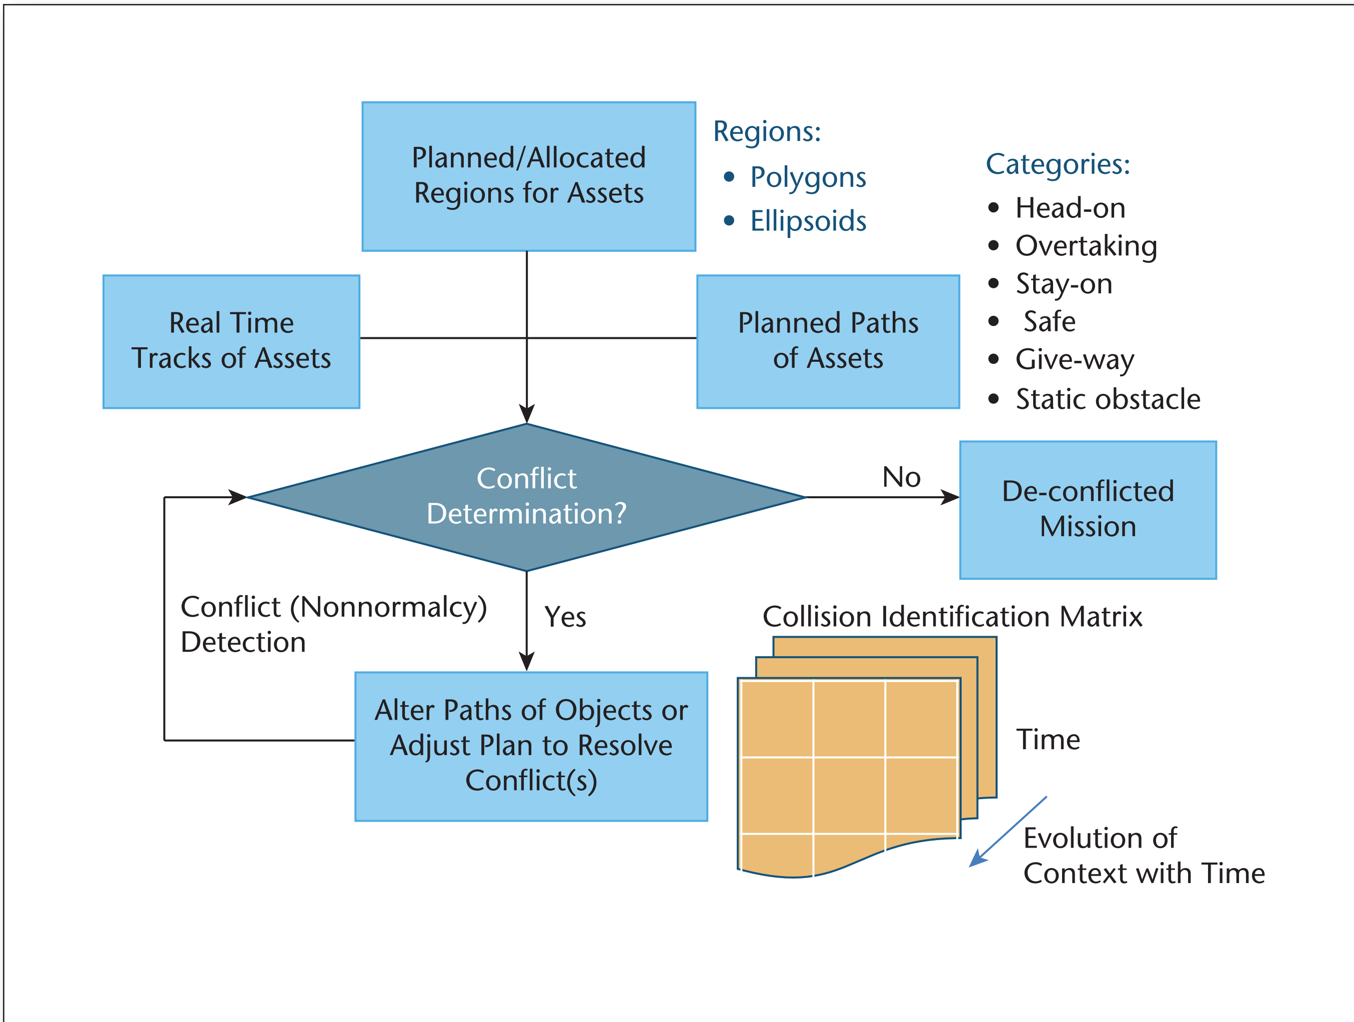

Figure 6. Conflict Identification.

and d). With severe weather degradation, proactive COAST assimilates this change in context and provides modified search boxes (that is, corrective measures to overcome a nonnormal scenario), reduced in size and shifted (change in search task). Because unfavorable weather conditions adversely affect asset performance, additional surveillance may be required (that is, a change in subgoals or mission) to concentrate on particular regions before routing interdiction assets to that location. This example illustrates how the change in environmental context ripples through the MEAT-H parameters and results in proactive COA recommendations by COAST.

Asset Package Selection and Planning for Maritime Operations Center

The dynamic resource management in maritime operations centers emphasizes standardized processes and methods, centralized assessments and guidance, networked distributed planning, and decentralized execution of missions across operations. A typical operational-level planning process in an abstract maritime operations center includes several intelligent entities. Future planners (FPs) collaboratively convert higher-level mission goals or a commander's intent into COAs for missions. Each COA decomposes mission goals into a graph of subgoals or specific tasks to achieve the goals and includes estimated requirements and available resources to accomplish every task. Each subgoal in the COA is represented as a task graph with branches and sequel options. Prior intelligence, historical and forecasted weather patterns, and logistics play key roles in developing future plans. Future operations (FOPS) allocate assets to tasks based on FP-specified COAs. This allocation is made over a moving time horizon (typically a 3-day horizon T, $T + 1$, $T + 2$, where T is the current day), taking into account dynamically evolving

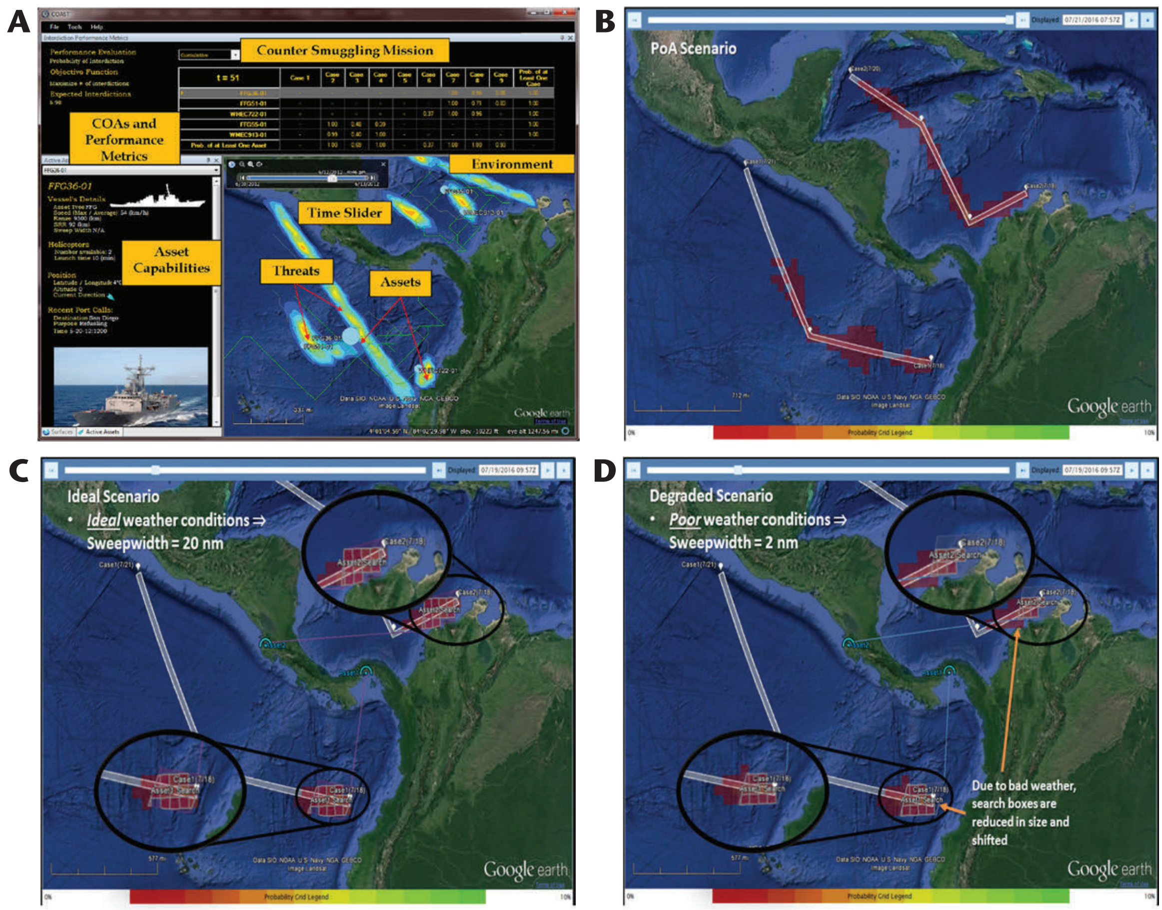

Figure 7. COAST.

(a) Original cases showing temporal evolution (indicated by color gradients) of the probability of smuggler activity in the eastern Pacific Ocean and the Caribbean Sea. (b) Target corridors (white rectangles) and uncertainty associated in red. (c) Ideal weather scenario with sweep width = 20 nm. (d) Degraded weather conditions affect the performance of sensors and sweep width = 2 nm (Mishra et al. 2017b).

intelligence, logistics, and weather information from reachback cells. Current operations (COPS) monitor the ongoing activities on day *T* and provide feedback to FOPS and FP in situation reports on emerging tasks and requirements, task outcomes, changes in task requirements, asset (and network) status, and evolving intelligence. Multilevel planning agents provide information and decision support to help FOPS planners evaluate and optimize asset-to-task allocations at several levels. At the execution level, agents suggest different supporting-supported options across a number of interdependent tasks in competing task graphs (representing different missions), taking into account uncertainty in weather forecasts, intelligence, and asset and network status. At platform and warfare area levels, agents optimize subplatform-to-task allocations. Other agents compute mission context–dependent values of information and decisions and manage the flow of information among decision makers. Details may be found in Han et al. (2014) and references therein.

Conclusion and Future Work

In this article, we demonstrated the use of IoT and AI for maritime decision making by identifying underlying contexts where different planning elements communicate to develop relevant courses of actions for human-machine systems. We discussed the algorithms for representation, diagnosis, and prognosis of context and methods to effectively communicate it across the data-to-decision process. We validated our algorithms via the development of proactive decision support tools across maritime operations, including (1) a multiobjective robust and adaptive optimization algorithms in the operational software tool, TMPLAR; (2) waterspace interference identification algorithms embedded in the operational software tool, CONFIDENT; (3) dynamic allocation of surveillance and interdiction assets for countersmuggling operations via COAST; (4) asset package selection and planning for MOC planning; and (5) machine learning and statistical hypothesis testing algorithms to infer cognitive context in digit recall, sequential letter recall, and arithmetic tasks using eye tracking data (SCOUT; Mishra et al. [2018]).

Our future research directions include practical interactive multiobjective optimization algorithms for resource allocation (for surveillance and execution) and asset routing, informed by context (mission, environment, asset, threat, human cognition) and data, featuring adaptive search interfaces (for example, scatter, gather for exploratory search tasks, baseline web search for lookup-type query tasks), Q-learning, multigrid methods, feature-based aggregation, rollout, deep reinforcement learning and approximate policy iteration, and modeling and analysis of cognitive context change detection in sequential tasks. We plan to use a unified graph-theoretical framework bringing together concepts from variational free energy optimization in thermodynamics and information theory; approximate dynamic programming from operations research and stochastic control; active inference-based perception and action selection from neuroscience; graphical model inference and bounded rationality from probabilistic inference and cognitive science; and Feynman–Kac path costs in physics to mathematically represent, evaluate, and design complex hybrid team structures.

Acknowledgments

The authors are thankful to Dr. William Lawless for his valuable time and excellent suggestions. The authors would like to thank US Office of Naval Research and Naval Research Laboratory for supporting this research.

Notes

1. www.docker.com.

2. Sweep width of any sensor is the width a definite range sensor has to sweep 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.

References

Abowd, G. D.; Dey, A. K.; Brown, P. J.; Davies, N.; Smith, M.; and Steggles, P. 2001. Towards a Better Understanding of Context and Context-Awareness. In *International Symposium on Handheld and Ubiquitous Computing*, 304–7. Berlin: Springer. doi.org/10.1007/3-540-48157-5_29

Alvarez, L.; Pérez-Peña, R.; and Robles, F. 2015. El Faro, Missing Ship, Has No Sign of Survivors. *New York Times* (October 5). www.nytimes.com/2015/10/06/us/el-faro-missing-ship-hurricane-joaquin.html

Bienkowski, A.; Zhang, L.; Sidoti, D.; and Pattipati, K. R. 2018. Path Planning in an Uncertain Environment Using Approximate Dynamic Programming Methods. In *21st International Conference on Information Fusion*. Piscataway, NJ: IEEE. doi.org/10.23919/ICIF.2018.8455762

Bonczek, R. H.; Holsapple, C. W.; and Whinston, A. B. 2014. *Foundations of Decision Support Systems*. Boston: Academic Press.

Chapanis, A. 1965. On the Allocation of Functions between Men and Machines. *Occupational Psychology* 39(1): 1–11.

Cummings, M. M. 2014. Man Versus Machine or Man+ Machine? *IEEE Intelligent Systems* 29(5): 62–9. doi.org/10.1109/MIS.2014.87

Fitts, P. M.; Viteles, M.; Barr, N.; Brimhall, D.; Finch, G.; Gardner, E.; Grether, W.; Kellum, W.; and Stevens, S. 1951. *Human Engineering for an Effective Air-Navigation and Traffic-Control System*. Oxford, UK: National Research Council.

Friston, K. 2010. The Free-Energy Principle: A Unified Brain Theory? *Nature Reviews. Neuroscience* 11(2): 127–38. doi.org/10.1038/nrn2787

George, D.; Lehrach, W.; Kansky, K.; La'Zaro-Gredilla, M.; Laan, C.; Marthi, B.; Lou, X. et al. 2017. A Generative Vision Model That Trains with High Data Efficiency and Breaks Text-Based CAPTCHAs. *Science* 358(6368). doi.org/10.1126/science.aag2612

Han, X.; Mishra, M.; Mandal, S.; Bui, H.; Ayala, D. F. M.; Sidoti, D.; Pattipati, K. R.; and Kleinman, D. L. 2014. Optimization-Based Decision Support Software for a Team-in-the-Loop Experiment: Multilevel Asset Allocation. *IEEE Transactions on Systems, Man, and Cybernetics. Systems* 44(8): 1098–112. doi.org/10.1109/TSMC.2013.2295360

Hansen, J. A.; Hodyss, D.; Bishop, C. H.; and Campbell, W. 2012. Coupled METOC/ INTEL Risk Assessment. U.S. Patent 13,272,272, filed April 19, 2012.

Horvitz, E. 1999. Principles of Mixed-Initiative User Interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 159–66. New York: ACM.

Kurien, J.; Nayak, P. P.; and Smith, D. E. 2002. Fragment-Based Conformant Planning. In *Proceedings of the Sixth International Conference on Artificial Intelligence Planning and Scheduling*, 153–62. Palo Alto, CA: AAAI Press.

Levchuk, G.; Pattipati, K. R.; Serfaty, D.; Fouse, A.; and McCormack, R. 2018. Active Inference in Multiagent Systems: Context-Driven Collaboration and Decentralized. Purpose-Driven Team Adaptation. Paper presented at 2018 AAAI Spring Symposium Series.Stanford, CA, March 26–28.

Levine, S. 2018. Reinforcement Learning and Control as Probabilistic Inference: Tutorial and Review. arXiv CoRR abstract:1805.00909. Ithaca, NY: Cornell University Library.

Mannaru, P.; Balasingam, B.; Pattipati, K.; Sibley, C.; and Coyne, J. T. 2017. Performance Evaluation of the Gazepoint gp3 Eye Tracking Device Based on Pupil Dilation. In *International Conference on Augmented Cognition*, 166–75. Berlin: Springer. doi.org/10.1007/978-3-319-58628-1_14

McMenemy, D.; Sidoti, D.; Palmieri, F.; and Pattipati, K. Forthcoming. A Fast and Efficient Conflict Detection Method for Ellipsoidal Safety Regions. *IEEE Transactions on Aerospace and Electronic Systems*.

Mishra, M.; An, W.; Sidoti, D.; Han, X.; Fernando Martínez Ayala, D.; Hansen, J. A.; Pattipati, K. R.; and Kleinman, D. L. Forthcoming. Context-Aware Decision Support for Anti-Submarine Warfare Mission Planning within a Dynamic Environment. *IEEE Transactions on Systems, Man, and Cybernetics. Systems*. doi.org/10.1109/TSMC.2017.2731957

Mishra, M.; Mannaru, P.; Sidoti, D.; Bienkowski, A.; Zhang, L.; and Pattipati, K. R. 2018. Context-driven proactive decision support for hybrid teams: supplementary material. Available at: https://www.researchgate.net/publication/327050002_Context-driven_Proactive_Decision_Support_for_Hybrid_Teams_Supplementary_Material_HUMAN-MACHINE_ADAPTIVE_TASK_ALLOCATION_FORMULATION

Mishra, M.; Han, X.; Sidoti, D.; Ayala, D. F.; An, W.; Kleinman, D. L.; and Pattipati, K. R. 2014. Multi-Objective Coordinated Path Planning for a Team of UAVs in a Dynamic Environment. Paper presented at 19th International Command and Control Research and Technology Symposium. Alexandria, VA, June 17–19.

Mishra, M.; Sidoti, D.; Avvari, G. V.; Mannaru, P.; Fernando Martinez Ayala, D.; and Pattipati, K. R. 2017a. Context-Driven Proactive Decision Support: Challenges and Applications. In *Computational Context: Why It's Important, What It Means, and Can It Be Computed? Papers from the 2017 AAAI Spring Symposium*. Technical Report SS-17-03. Palo Alto, CA: AAAI Press. aaai.org/Library/Symposia/Spring/ss17-03.php

Mishra, M.; Sidoti, D.; Avvari, G. V.; Mannaru, P.; Ayala, D. F. M.; Pattipati, K. R.; and Kleinman, D. L. 2017b. A Context-Driven Framework for Proactive Decision Support with Applications. *IEEE Access : Practical Innovations, Open Solutions* 5: 12475–95. doi.org/10.1109/ACCESS.2017.2707091

Muir, B. M. 1987. Trust Between Humans and Machines, and the Design of Decision Aids. *International Journal of Man-Machine Studies* 27(5-6): 527–39. doi.org/10.1016/S0020-7373(87)80013-5

Parasuraman, R.; Sheridan, T. B.; and Wickens, C. D. 2000. A Model for Types and Levels of Human Interaction with Automation. *IEEE Transactions on Systems, Man, and Cybernetics. Part A, Systems and Humans* 30(3): 286–97. doi.org/10.1109/3468.844354

Perera, C.; Zaslavsky, A.; Christen, P.; and Georgakopoulos, D. 2014. Context Aware Computing for the Internet of Things: A Survey. *IEEE Communications Surveys and Tutorials* 16(1): 414–54. doi.org/10.1109/SURV.2013.042313.00197

Pomerol, J.-C. 1997. Artificial Intelligence and Human Decision Making. *European Journal of Operational Research* 99(1): 3–25. doi.org/10.1016/S0377-2217(96)00378-5

Powell, W. B. 2007. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. New York: John Wiley & Sons. doi.org/10.1002/9780470182963

Rasmussen, J. 1983. Skills, Rules, and Knowledge; Signals, Signs, and Symbols, and Other Distinctions in Human Performance Models. *IEEE Transactions on Systems, Man, and Cybernetics* SMC-13(3): 257–66. doi.org/10.1109/TSMC.1983.6313160

Sidoti, D. 2016. Why Context and Context-Driven Decision Support Matters. *IEEE Aerospace and Electronic Systems Magazine* 31(11): 50–2. doi.org/10.1109/MAES.2016.150251

Sidoti, D.; Avvari, G. V.; Mishra, M.; Zhang, L.; Nadella, B. K.; Peak, J. E.; Hansen, J. A.; and Pattipati, K. R. 2016. A Multiobjective Path-Planning Algorithm with Time Windows for Asset Routing in a Dynamic Weather-Impacted Environment. *IEEE Transactions on Systems, Man, and Cybernetics. Systems* 47 (12): 3256–71. doi.org//10.1109/TSMC.2016.2573271

Sidoti, D.; Han, X.; Zhang, L.; Avvari, G. V.; Ayala, D. F. M.; Mishra, M.; Sankavaram, M. S.; Kellmeyer, D. L.; Hansen, J. A.; and Pattipati, K. R. Forthcoming. Context-Aware Dynamic Asset Allocation for Maritime Interdiction Operations. *IEEE Transactions on Systems, Man, and Cybernetics. Systems* doi.org/10.1109/TSMC.2017.2767568

Smirnov, A. 2006. *Context-Driven Decision Making in Network-Centric Operations: Agent-Based Intelligent Support*. Fort Belvoir, VA: Defense Technical Information Center.

Sutton, R. S., and Barto, A. G. 1998. *Reinforcement Learning: An Introduction*. Vol. 1. Cambridge, MA: MIT Press.

Tracy, P. 2016. Industrial Internet of Things Maritime Use Cases. enterpriseiotinsights.com/20160727/Channels/Use-Cases/Maritime-Industrial-Internet-of-Things-tag31-tag99.

Wakabayashi, D. 2018. Self-Driving Uber Car Kills Pedestrian in Arizona, Where Robots Roam. *New York Times* (March 19). www.nytimes.com/2018/03/19/technology/uber-driverless-fatality.html.

Yedidia, J. S.; Freeman, W. T.; and Weiss, Y. 2005. Constructing Free-Energy Approximations and Generalized Belief Propagation Algorithms. *IEEE Transactions on Information Theory* 51(7): 2282–312. doi.org/10.1109/TIT.2005.850085

Zhang, S.; Pattipati, K. R.; Hu, Z.; Wen, X.; and Sankavaram, C. 2013. Dynamic Coupled Fault Diagnosis with Propagation and Observation Delays. *IEEE Transactions on Systems, Man, and Cybernetics. Systems* 43(6): 1424–39. doi.org/10.1109/TSMC.2013.2244209

Manisha Mishra received her BS in electrical engineering from Punjab Technical University 2004; MS in electrical and computer engineering from the University Hawaii 2008; and her PhD in electrical and computer engineering from the University Connecticut 2017. She was a software engineer with Infosys Technologies, India 2005. She worked as an ORISE post-graduate researcher with US Army Engineering Research and Development Center-Cold Region Research Labs 2017. She is an algorithm engineer with Aptiv Corporation. Her research includes modeling dynamic and uncertain environments for asset allocation and path planning, context aware decision support systems, risk analysis, system diagnosis and prognosis, and optimization-based techniques for mission planning and coordination.

Pujitha Mannaru received a BE degree in electronics and communications engineering from PES Institute of Technology - Bangalore South Campus, affiliated with Visvesvaraya Technological University, India, in 2013. She is currently pursuing her PhD in electrical and computer engineering at the University of Connecticut. Her research interests are in the areas of applications of engineering psychology, signal processing, and machine learning to proactive decision support tools and human-machine systems.

David Sidoti received BS, MS, and PhD degrees in electrical and computer engineering from the University of Connecticut, Storrs in 2011, 2016, and 2018, respectively. He is currently a computer scientist at the U.S. Naval

Research Laboratory – Marine Meteorology Division. He was the corecipient of the Tammy Blair award for best student paper at FUSION 2016. His current interests include multiobjective algorithms for dynamic scheduling and resource management in weather-impacted environments. His additional research foci include information valuation, combining optimization techniques with deep learning, and sailing vessel routing.

Adam Bienkowski received BS and MEng degrees in electrical and computer engineering from the University of Connecticut, Storrs, CT in 2013 and 2017, respectively, and is currently pursuing the PhD degree with the Department of Electrical and Computer Engineering at the University of Connecticut under the advisement of Dr. K. R. Pattipati. He was an electrical engineer at General Dynamics Electric Boat, Groton, CT from 2013 to 2017. His current research interests include modeling dynamic and uncertain environments for asset allocation and path planning, context aware decision support systems, and optimization and machine learning-based techniques for mission planning and coordination.

Lingyi Zhang received a BS degree in electrical and computer engineering from the University of Connecticut, Storrs, CT, USA, in 2014, where she is currently pursuing a PhD degree. Her current research interests include modeling dynamic and uncertain environments for asset allocation and path planning, context-aware decision support systems, and optimization-based techniques for mission planning and coordination. Zhang was a corecipient of the Tammy Blair Award for Best Student Paper at FUSION 2016.

Krishna R. Pattipati is the Board of Trustees Distinguished Professor and the UTC Chair Professor of Systems Engineering in the department of Electrical and Computer Engineering at the University of Connecticut, Storrs, CT, USA. Pattipati's research activities are in the areas of proactive decision support, autonomy, and optimization-based learning and inference. A common theme among these applications is that they are characterized by a great deal of uncertainty, complexity, and computational intractability. He is a Life Fellow of IEEE.