Towards a Context-Aware Proactive Decision Support Framework

Benjamin B. Newsom, Jr[1]
Next Century Corporation
Columbia, Maryland, USA

Ranjeev Mittu[2]
Naval Research Laboratory
Washington, DC, USA

Ciara Sibley[2]
Naval Research Laboratory
Washington, DC, USA

Myriam Abramson[2]
Naval Research Laboratory
Washington, DC, USA

[1]ben.newsom@nextcentury.com, [2]{firstname.lastname@nrl.navy.mil}

Abstract—**The problem of automatically recognizing a user's operational context, the implications of its shifting properties, and reacting in a dynamic manner is at the core of mission intelligence and decision making. Environments such as the OZONE Widget Framework[1] provide the foundation for capturing the objectives, actions and activities of the mission analyst and decision maker. By utilizing a "context container" that envelops an OZONE Application, we hypothesize that *action* and *intent* can be used to characterize user context with respect to operational modality (strategic, tactical, opportunistic, or random). As the analyst moves from one operational modality to another, we propose that information visualization techniques should adapt and present data and analysis pertinent to the new modality and to the trend of the shift. As a system captures the analyst's actions and decisions in response to the new visualizations, the context container has an opportunity to assess the analyst's perception of the information value, risk, uncertainty, prioritization, projection and insight with respect to the current context stage. This paper will describe a conceptual architecture for an adaptive work environment for inferring user behavior and interaction within the OZONE framework, in order to provide the decision-maker with context relevant information.**

Keywords—context-driven; decision-making; dynamic modeling; operational modality; temporal reasoning

I. INTRODUCTION

Today's warfighters operate in a highly dynamic world with a high degree of uncertainty, compounded by competing demands. Timely and effective decision making in this environment is increasingly challenging. The phrase "*too much data – not enough information*" is a common complaint in most Naval operational domains. Finding and integrating decision-relevant information (vice simply data) is difficult. Mission and task context is often absent (at least in computable and accessible forms), or sparsely/poorly represented in most information systems. This limitation requires decision makers to mentally reconstruct or infer contextually relevant information through laborious and error-prone internal processes as they attempt to comprehend and act on data. Furthermore, decision makers may need to multi-task among competing and often conflicting mission objectives, further complicating the management of information and decision making. Clearly, there is a need for advanced mechanisms for the timely extraction and presentation of data that has value and relevance to decisions for a given *context*.

To put the issue of context in perspective, consider the fact that nearly all national defense missions involve Decision Support Systems (DSS)—systems that aim to decrease the cycle time from the gathering of data to some operational decision. The proliferation of sensors and large data sets are overwhelming DSS's, as they lack the tools to efficiently process, store, analyze, and retrieve vast amounts of data. Additionally, these systems are relatively immature in helping users recognize and understand important context (i.e. cues). The next generation systems must leverage predictive models to enable Proactive Decision Support (PDS). These systems will need to understand and adapt to user context (missions, goals, tasks). By aligning the data with the user in the appropriate context, we hypothesize that more relevant information can be provided to the user i.e., likely to be of higher value for decision making. The key challenges, therefore, are to not only model the user's decision-making context, but to recognize when such context has shifted. With regard to Figure 1, we hypothesize that concepts associated with PDS closely align with Prescriptive Analytics (i.e., understanding and modeling decision trajectories and the relevant information necessary for those decisions).

Descriptive Analytics	Predictive Analytics	Prescriptive Analytics
Answers the question, "*What happened?*" Examines data to identify trends and patterns.	Answers the question, "*What might happen in the future?*" Uses Predictive Models to Forecast Future.	Answers the question, "*What is the best decision to take given the predicted future?*"

Figure 1: Comparison of different forms of Analytics

The problem of automatically recognizing / inferring user context, understanding the implications of its shifting properties, and reacting in a dynamic manner is at the core of mission intelligence and decision making. An environment such as the OZONE Widget Framework provides the foundation for capturing the objectives, actions and activities of the mission analyst/decision maker. By utilizing a "context container" that envelops an OZONE Application, we can capture both action and intent which allows us to characterize this context with respect to its operational modality (strategic, tactical, opportunistic, or random) – Figure 2 (*Visual Analytics* representation).
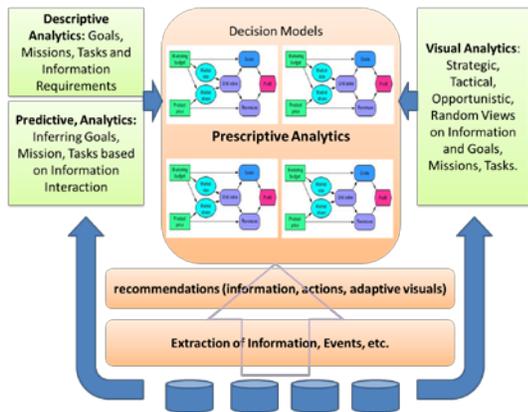

Figure 2: Context understanding in relation to Analytics

Context is fluid over time, and the relative mix of strategic vs. tactical vs. opportunistic actions or activities is also changing. Knowing the time frame and distribution of activities gives us insight into the analyst's changing operational modality. A temporal storage approach, such as a Context-Aware Memory Structure (CAMS), provides the basis for comparison of the "current" decision stage against prior stages and is used to predict phase shift.

Methods for understanding user context can be found in logic-based or probabilistic Artificial Intelligence (AI) approaches under Predictive Analytic Methods, or through more traditional methods based on Descriptive Analytics. Using a Descriptive Analytics approach, models can conceivably be developed that map missions, goals and tasks to information requirements in order to represent "decision context". With regard to deriving context within the Predictive and Visual analytics models, the challenging questions become: Can a user's decision context be modeled, based upon, information seeking, interaction, or analysis patterns [1]? What research can be leveraged from the AI community (plan recognition) to infer which decision context (model) is active? Can we reason about which decision context (model) should be active? What similarity metrics enable the selection of the appropriate model for a given context? Can we recognize context shift based on work that has been done in the Machine Learning community with "concept drift", and how well does this approach adapt to noisy data? The emphasis for the paper will be

on the *Visual Analytics* representation for understanding context, but the questions span across *the Predictive Analytics* representation as well.

In Section II, we provide a notional operational example to guide the framework discussion. In Section III, we describe the APTO system architecture. In Section IV, we briefly describe the idea of Context Container for the APTO framework. In Section V and VI we describe the Context Aware Memory Manager and context shift recognition. In Sections VII, VIII, IX, and X we discuss the adaptive visualization informed through the APTO architecture, event, activity and workflow manager, respectively.

II. NOTIONAL OPERATIONAL EXAMPLE

Consider the scenario of the intelligence analyst on a *24x7* watch floor (Figure 3). As the analyst moves from one operational modality to another, the information visualization techniques should adapt and present data and analysis pertinent to the new modality and to the trend of the shift. If we can capture the analyst's actions and decisions in response to the new visualizations, the context container may be able to infer the analyst's perception of the information value, risk, uncertainty, prioritization, projection and insight. This information, in combination with the ability to infer the user's current context stage would provide the ability for DSS's to pre-stage information that is tailored to the user's current needs and preferences along a decision trajectory.

Each watch floor is configured and organized to address their unique and specialty mission and intelligence requirements. As such, any solution proposed must be able to adapt and conform to the specific needs of the watch. In Figure 3, we show an example set of watch floor responsibilities with the proposed solution focusing on Analyst activities (3), Cell activities (4), and Watch Officer activities (6).
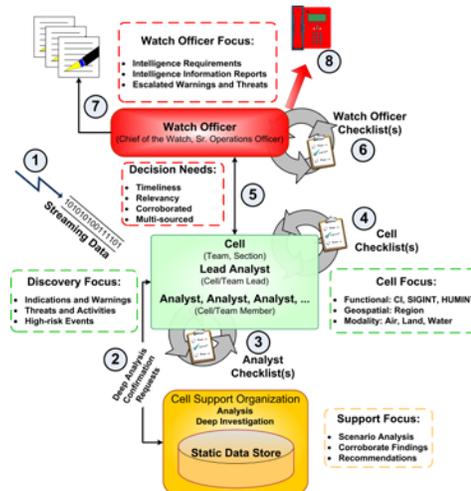

Figure 3: Example Watch Floor Scenario

In general, a watch floor is organized around Cells of responsibility. A Cell (also known as a *Team* or *Section*)

may have only one Analyst with a singular focus, or it may be multiple Analysts with a Lead Analyst (also known as the Cell or Team Lead). A Cell is monitoring and accumulating streaming data (1) to discover indications and warnings about threats and high-risk events in their scope of consideration. Timeliness of analysis and interpretation is critical. The Cell may have a support organization that can perform deep analysis (2) and confirm an Analyst's or Cell's findings. For often-detected indications, the Analyst will have a set of standard operating procedures or checklist (3) of activities they need to perform to reach the decision to escalate the detected event to the next level. In a multi-person Cell, the next level may be the Cell Officer who has their own set of standard operating procedures or checklist (4) of activities that need to be performed to escalate out of the Cell (5).

An event (threat or warning) escalated out of the Cell (5) goes to the Watch Officer who is accumulating information and comparing escalated events to their Intelligence Requirements. Like the Analyst and the Cell, the Watch Officer has a set of standard operating procedures or checklist (6) of activities to perform in response to the combination of escalated events that they are receiving from all of the Cells on the watch floor. The Watch Officer makes the trade-off decisions to only track and log (7) the events (threats) or escalate identified, confirmed, credible threats (8) to the next level.

The watch floor situation has intense analytical problems requiring timely analyses and/or responses. Analytical problems are often sensitive and associated with high stakes for success or failure. In many analytical sub-domains, the objectives of the analysis can be open and shifting, and analysts must sometimes determine for themselves the goals and priorities of their data collection or research. The proposed framework identifies the context in which the events and activities are occurring and provides situational awareness and accuracy up and through the chain of decision makers.

The proposed system architecture should extend and enhance existing mission solutions to include PDS focusing on context shift recognition and staging of the information (or combinations of information) the analyst requires in making the "next" decision. Along with determining the information to be staged, the adaptive work environment needs to react to the context shift and determine the appropriate stage-related information visualization techniques.

To accomplish the objective of inferring a user's context and recognizing context shifts, there are three broad areas of required innovation:

- Capturing context actions and events through normal analyst interaction with OZONE Framework applications.

- Characterizing the user's actions and events along their operational modality (i.e., strategic, tactical,

random discovery, and opportunistic discovery), their temporal relationship, and situational objectives.

- Recognizing the change or shift in context through the development of Context Shift Models and predictive analysis.

III. APTO SYSTEM

A. Long Term Goal

In order to create a context-aware adaptive work environment, specific elements such as the memory components, the context manager, and the Activity Manager are necessary for recognizing context and context shift. APTO (Latin for adapt) is a conceptual architecture, shown in Figure 4, that depicts a context-aware environment within the OZONE Widget Framework.
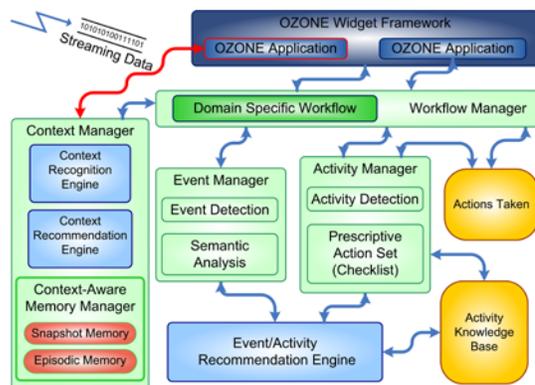

Figure 4: Conceptual APTO Architecture

B. Technical Approach

The premise of our approach is that the combination of an intelligence analyst's OZONE Application (sometimes referred to as widgets) usage pattern and the information being visualized (and how it is visualized) can be used as an indicator of the analyst's context mode. The analyst is viewing all of the situation characteristics through a particular lens searching for strategic insights, tactical clues, opportunistic indicators or the random-scramble searching for the information nugget that connects decision streams together. Through adaptation and innovative extensions to the OZONE Widget Framework, it will be possible to capture traces of user interactions with the widgets, as well as interactions between widgets. We believe this situational capture of the decision making process will form distinctive, predictable patterns of behavior corresponding to the analyst's intent, information value, and prioritization.

IV. CONTEXT CONTAINER FOR OZONE APPS

The concept of a "context container" for OZONE Apps does not exist in the current OZONE Widget Framework. In the overall architecture, it is part of the

interaction between the user experience or presentation layer and the Context Manager. We believe that we can define and create a container or software envelope that would "wrap" an OZONE application (a collection of one or more widgets and data sources) and automatically capture both what the decision was and an indication of why (operational modality) the decision or choice was made. This collection of activity, interaction and/or decisions represents a context vector that would be stored in the Context-Aware Memory Manager.

V. CONTEXT-AWARE MEMORY MANAGER

To model the analyst's context, a learning context memory model (a Context-Aware Memory Manager – CAMS) [2] could be constructed. This model would capture the OWF widget interactions and process them to construct a context memory reflecting the user's regular activity. The concept of a Context-Aware Memory Manager that interacts with the OZONE Widget Framework does not currently exist.

Context memory is a mechanism for retaining and recalling interesting and relevant past experiences or actions [3]. We believe that an analyst's context consists of a striation or mix of strategic, tactical, opportunistic and random actions. In each layer there are a collection of short-term or "snapshot" memories and long-term or "episodic" memories.

A. Snapshot Memory

The snapshot memory (context working memory) processes and stores context attributes from context input vectors. Attributes are stored in Artificial Recognition Balls [4] (ARBs), which describe a certain region around the context attribute —in the case of OZONE Apps it would be the context container— and enables CAMS to perform data compression by eliminating the need for repetition. For example, a particular type of action can be represented by a single ARB instead of all individual actions that occur within the container; every ARB has a resource level R associated with it, being an indicator for how frequently it recognizes context attributes. The algorithm used in CAMS is based on the principles of unsupervised and reinforcement learning. Unsupervised learning allows us to construct a system which can cluster input data without any prior knowledge about the structure of every class. Reinforcement learning requires feedback from a trainer. However, an explicit trainer is not desirable in most context-aware systems, therefore an ARB receives positive feedback (stimulation) when context attributes fall within a certain distance from the center, resulting in an increase in its resource level. Negative feedback is introduced by the notion of 'forgetting', which gradually decays all resource levels. For example, actions a user performs less often have their resource level reduced by a decay factor, but every re-occurrence stimulates it again, which enables these actions to remain in memory.

B. Episodic Memory

To capture a significant part of human activity the connections between consecutive events or actions are essential. The snapshot memory is able to capture every individual action, but not the set of actions that comprise a specific decision. As the user is most likely to login/logout, start up an application, etc., those actions have a higher resource level R. Once R reaches a predefined level, the oft repeated actions are passed from the Snapshot Memory to the Episodic Memory, which captures all individual attribute values between them. The Context Memory Manager component regulates the division of the memory mechanism into Snapshot and Episodic Memory. This division is essential for keeping the complexity of the search space at a manageable level. Without this division all attributes and connections between them would have to be stored in a directed graph in order to detect and capture meaningful consecutive events — which would result in an *NP* complete search problem. Instead, only the attribute vectors between ARBs with a high resource level need to be stored; after the validation of an episode this is reduced to storing only references to ARBs recognizing the attributes in these vectors. The ARBs with a high resource level R passed on from the Snapshot Memory are stored in a cache structure.

Initially, the user will be asked to name and validate a new or preliminary episode, bridging the gap between the data representation within CAMS and the real world meaning. An episode is an ordered 3-tuple containing a start ARB, an end ARB and an ordered list of all context vectors encountered. Ideally, in order for the proposed system to diffuse into every day environments, APTO could learn from the human-assisted validation and move towards automatic recommendations for naming and validation. Only frequently occurring episodes would be presented.

VI. CONTEXT SHIFT MODEL AND SHIFT RECOGNITION

We believe that we can create a network model of the ordered 3-tuple activities that represent each of a context mode's three stages: entering a mode, "*in-the-flow*" of a mode, and exiting a mode based on user interaction patterns. These context mode stage models can be compared to a dynamic modeling of the analyst's real-time activities for detecting shifts and flows of focus. Each mode stage (entering, in-flow, leaving) is a combination or mix of the operational modalities (strategic, tactical, opportunistic, or random) within a particular time frame.

This mix is constantly changing as new information is being presented to the analyst. This combination of actions (e.g., 80% strategic, 12% tactical, 6% opportunistic and 2% random) collected from the analyst's interaction with APTO, will provide the context profile for that analyst at that given time. As

they interact with APTO, their profile trend changes, thus their context and items of interest change.

In particular, the user experience activity of "zooming in" on the temporal aspect of streaming data typically characterizes a tactical desire to narrow the focus for an immediate decision. Typically, this behavior is followed by a "zooming out" to take a more strategic view of the information looking for particular clusters of relevant events or activities. Although this is typical, not all analysts operate in the same manner. Our proposed approach is to accommodate an individualized recognition of pattern and transition indicators [5]. By capturing usage patterns and successful episodes on an individualized basis, the system will be able to adapt its shift recognition to the specific analyst. Over time, the patterns accumulated could become the basis for identification of a best practice approach for often repeated situations.

VII. CONTEXT SHIFT-AWARE STAGING AND VISUALIZATION

Our "context shift" goal is to deliver an individual-focused, context-aware component that can feed its analysis and recognition of transition stages to our context-aware components so that they can anticipate and pre-stage data and recommendations. The analyst's "*view of the world*" should adapt to the individual's operational modality (strategic, tactical, opportunistic, or random). This includes recognizing the data sources, widgets and visualization techniques that are applicable to the particular mode. This identification process will rely heavily upon the context container that encompasses and defines the operational characteristics of the OZONE App.

VIII. EVENT MANAGER

The basis for the Event Manager comes from the Event Representation and Structuring of Text (EVEREST) project, sponsored by the Office of Naval Research. It is an SBIR initiative that has developed text analytic technology that crosses the semantic gap into the area of event recognition and representation. The EVEREST system searches for mappings to a semantic event model, interactively suggesting evidence for the occurrence of whole or partial events for human analysis and reporting. The semantic targeting approach extends the ideas of Open Information Extraction [6], Event Web [7], Semantic Web [8], and the OZONE Widget Framework. EVEREST's event-centric approach is critical for generating narratives that confer meaning upon large, complex, uncertain, and incomplete data sets.

A. Event Detection

The event detection component is based on an Open Information Extraction (Open IE) [9] approach. Open IE systems distill huge volumes of text into a long list of tuples (two entities and one relation that binds them) without asking a human for examples of those relations first. We consider each entity→relationship→entity tuple to be an event assertion. The extractions of assertions from the text are entirely lexical in nature. The assertion extraction utilizes Stanford's core NLP libraries and makes use of a *part-of-speech* tagger (annotator) and noun phrase "chunker." To locate the word in the vicinity of the two nouns (or noun phrases) that mostly likely intended to express their relationship, the detection algorithm employs a technique known as conditional random fields. In essence, this is a statistical model that is sensitive to its lexical context.

B. Prescriptive Event Recognition

The Prescriptive Event Recognition component comprises an event semantic model (metadata and list of assertions) and event inference engine that compares predetermined Target Event models with Reports (detected metadata and list of assertions) in the input stream. The event semantic model is based on Wasterman and Jain [10].

The event inference engine is a mixed-initiative application, i.e., one with a human in the loop, which compares extracted assertions against a prescribed model using a rules-over-graphs approach. The key idea is that many inferencing algorithms used by logic-based AI systems can be heuristically approximated by a much simpler and more efficient system based on graph-matching algorithms. The assertions associated with a Target Event are modeled as a graph of nodes and edges. The nodes are the entities of the tuple. The edges are the relationships between the entities. Similarly, the event assertions detected in the incoming data stream are modeled as a graph of nodes and edges. The graphs are compared for shape, structure, directionality of the edges, content (metadata) of the nodes, and content (metadata) of the edges. Each comparison is scored or ranked to determine how closely the detected event assertion matches the Target Event.

The Prescriptive Event Recognition component offers a list of assertions that are candidate matches for a Target Event. The initial list of candidate assertions are ranked by the inference engine based on its searches for class, instance, and relation isomorphisms between all of the assertions and its semantic event models; an assertion with a closer resemblance will find itself higher on the list. The informational value of the assertion—whether it would fill a central node or an outlier in the graph—will influence the rank as well. The user can decide to accept (or reject) the assertion after consulting his own knowledge, source documents, or other materials. This process could be utilized to fill in missing parts of a graph, which in turn could be utilized by the system to uncover new pieces of information, and this cycle would continue until a target concept has been proven.

IX. ACTIVITY MANAGER

The Activity Manager is focused on activities that are occurring inside the APTO architecture. It interacts with

OZONE Applications via the context container, with the Context Manager module, and the Actions Taken repository.

A. Action Detection

The Action Detection component interacts with OZONE Applications via the context container and the Actions Taken repository. It monitors all of the activities occurring within APTO and identifies actions of interest to the Domain Specific Workflow and routes these actions to the Prescriptive Action component.

B. Prescriptive Action Set

The Prescriptive Action component comprises an action semantic model (metadata and a list of assertions) and an activity inference engine that matches predetermined Action Sets (checklists) with Events (detected metadata and a list of assertions) and Actions Taken. Similar to the common event model proposed by Wasterman and Jain [10], the action semantic model contains temporal elements (the time horizon over which the action should occur), spatial elements (the geographic location where the action should occur), structural elements (the set of action assertions, process steps, or checklist items that need to occur), informational elements (the actor that should perform the action), and causal elements (the set of event assertions that caused this particular action model to be selected).

C. Suadeo Recommendation Engine

Suadeo (Latin for recommendation) is a prototype context-aware, model-driven, recommender system that utilizes "static" persistent data and streaming data as the basis for deriving its recommendations. The intent of the Suadeo prototype is to be a hybrid recommender system that is context-aware with the context model being defined along multiple dimensions such as person, place, time, and incident. The recommendation engine is driven from a graph-based analysis of the Actions Taken metadata and tuples. Although the description of the recommendation engine in the context of Figure 4 is to provide a predefined set of actions in the form of recommended checklists, in the more general setting the recommendations could be new information sources that might be relevant for a given decision.

One of the challenges with regard to the development of a recommendation engine is how the system should "understand" and adapt to the various biases inherent in the way humans explore their information environment? For example, information bias (the tendency to seek information even when it cannot affect action), confirmation bias (the tendency to search for or interpret information or memories in a way that confirms one's preconceptions) and anchoring (the tendency to rely too heavily, or anchor, on one trait or piece of information when making decisions) may be guiding the humans information seeking patterns. Any recommender system, through its ability to better manage and understand user-context and the decision making environment, should help overcome these limitations.

X. WORKFLOW MANAGER

Although the specific example of a 24x7 Watch Floor is used to describe the concepts of APTO, the intent of the architecture is to accommodate a broader class of problems. The general characteristics of these problems are that they have a high volume of streaming and static data that is composed of structured components and unstructured data (predominately text data). The unstructured data can be given structure in the form of an event assertion (a semantic tuple). From the combination of the original structured components and the discovered event assertions, events can be determined. Once an event (or set of events) is determined, a set of actions needed to respond to the event can be determined. In many, but not all, situations, it is desired that the system identify, track and remember the actions taken.

Depending upon the specific domain or scenario being addressed by APTO, only some of these process steps are required to reach the objective of having actionable information upon which to make a decision. To accommodate different workflows (or process steps), the APTO architecture is comprised of independent, re-usable modules whose interactions represent a workflow. Every module in the architecture reports what it has done to the Workflow Manager. For example, when a new event assertion is created, the Workflow Manager is notified. Based upon the notification received and the specific workflow that is being executed, the next process step is determined and executed. It is envisioned that there may be multiple concurrent workflows executing within APTO.

A. Domain Specific Workflow

A Domain Specific Workflow component defines how data (objects) flow through the APTO architecture, determines which Action Taken items are important, and which Action Taken items trigger new Activities (or Action Sets).

B. Actions Taken

The Actions Taken component contains all of the actions that have occurred within the APTO architecture. Similar to our Target Events and Reports, the Action Taken domain object is a collection of metadata and a list of assertions (tuples). Essentially, an Action Taken item is a realized instance of an action semantic model. Where the model in the Prescriptive Action Set identifies what "should" occur, the Action Taken object identifies what actually happened answering the "Who", "What", "When", "Where", and "Why" questions.

CONCLUSION

This paper has discussed a context aware Proactive Decision Support framework within the OZONE environment. Furthermore, several longer term challenges have been briefly described with regard to modeling decision context, metrics for recognizing operational context, and techniques for recognizing context shift. Additional research areas include:

- Adequately capturing users' information interaction (seeking) patterns (and subsequently user information biases)

- Reasoning about information seeking behaviors in order to infer decision making context; for example, the work being done by researchers within the Contextualized Attention Metadata community [11] and the Universal Interaction Context Ontology [12] might serve as a foundation

- Instantiating formal models of decision making based on information seeking behaviors

- Leveraging research from the AI community in plan recognition to infer which decision context (model) is active, and which decision model should be active

- Recognizing decision shift based on work that has been done in the Machine Learning community with "concept drift", and assessing how well this approach adapts to noisy data and learns over time

- Incorporating uncertainty and confidence metrics when fusing information and estimating information value in relation to decision utility

Elaborating further on the ideas presented in the paper, longer term research should be focused on the following:

Decision Models for goal-directed behavior: Instantiation of prescriptive models of decision making, which integrate information recommendation engines that are context-aware. Furthermore, techniques that can broker across, generalize, or aggregate, individual decision models would enable application in broader contexts such as group behavior. Supporting areas of research may include similarity metrics that enable the selection of the appropriate decision model for a given situation, and intuitive decision model visualizations.

Information Extraction and Valuation: Locating, assessing, and enabling, through utility-based exploitation, the integration of high-value information within the decision models, particularly in the big data realm is a research challenge due to the heterogeneous data environment. In addition, techniques that can effectively stage relevant information along the decision trajectory (while representing, reducing and/or conveying information uncertainty) would enable the wealth of unstructured data to be maximally harnessed.

Decision Assessment: Modeling decision "normalcy", in order to identify decision trajectories that might be considered outliers and detrimental to achieving successful outcomes in a given mission context would be areas for additional research. Furthermore, techniques that proactively induce the correct decision trajectory to achieve mission success are also necessary. Lastly, metrics for quantifying decision normalcy in a given context can be used to propose alternate sequences of decisions or induce the exact sequence of decisions. This would require the pre-staging of the appropriate information needed to support the evaluation of those decisions and would potentially improve the speed and accuracy of decision making.

Operator/Human Issues: Understanding, modeling and integrating the human decision making component as an integral part of the aforementioned areas is a novel areas of research. The challenges are to represent human decision-making behavior computationally, to mathematically capture the human assessment of information value, risk, uncertainty, prioritization, projection and insight; and computationally representing human foresight and intent.

REFERENCES

[1] Filip Radlinski, Martin Szummer, Nick Craswell. Inferring query intent from reformulations and clicks. • Proceedings of the 19th international conference on World wide web. Pages 1171-1172. ACM New York, NY.

[2] Mohr, P. H., Ryan, N., & Timmis, J. (2006). Capturing Regular Human Activity through a Learning Context Memory. In *Proceedings of the 3rd International Workshop of Modelling and Retrieval of Context (MRC 2006) in conjunction with AAAI-06* (p. 6).

[3] Mohr, P.; Timmis, J.; and Ryan, N. (2005), Immune inspired context memory. In *1st International Workshop on Exploiting Context Histories in Smart Environments*, 4.

[4] Neal, M. (2003), *Meta-stable Memory in an Artificial Immune Network, Proceedings of the 2nd International e-Conference on Artificial Immune Systems*, p. 229-241.

[5] Agrawal, Vikas, Heredero, Genoveva, Penmetsa, Harsha, Laha, Arijit, and Shastri, Lokendra. "Activity Context Aware Digital Workspaces and Consumer Playspaces: Manifesto and Architecture" AAAI Workshops (2012): n. pag. Web. 14 Aug. 2013

[6] Michele Banko, Michael J. Cafarella, Stephen Soderland, Matthew Broadhead, and Oren Etzioni. Open information extraction from the web. In IJCAI, pages 2670-2676, 2007.

[7] Ramesh Jain: EventWeb: Developing a Human-Centered Computing System. IEEE Computer 41(2): 42-50 (2008)

[8] Jim Hendler: Web 3.0 Emerging. IEEE Computer 42(1): 111-113 (2009).

[9] Anthony Fader, Stephen Soderland, Oren Etzioni: Identifying Relations for Open Information Extraction. EMNLP 2011: 1535-1545.

[10] Utz Westermann, Ramesh Jain: Toward a Common Event Model for Multimedia Applications. IEEE MultiMedia 14(1): 19-29 (2007).

[11] http://www.dlib.org/dlib/september07/wolpers/09wolpers.html, retrieved on 1 November 2013.

[12] A. Rath, D. Devaurs, and S. Lindstaedt. UICO: an ontology-based user interaction context model for automatic task detection on the computer desktop. In CIAO '09: Proceedings of the 1st Workshop on Context, Information and Ontologies, page 10, NewYork, NY, USA, 2009. ACM