SEM^2 Suite — Towards a Tool Suite for Supporting Knowledge Management in Situation Awareness Systems

Andrea Salfinger*, Daniel Neidhart‡, Werner Retschitzegger§, Wieland Schwinger¶
Department of Cooperative Information Systems
Johannes Kepler University Linz
4040 Linz, Austria
{andrea.salfinger*, daniel.neidhart‡, werner.retschitzegger§, wieland.schwinger¶}@cis.jku.at

Stefan Mitsch†
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA, USA
smitsch@cs.cmu.edu

***Abstract*—Human operators in today's control centers, such as air or road traffic control, need to monitor a plethora of information obtained from diverse sources. To support them in detecting critical situations within this information flood and taking timely actions, operators thus need adequate information fusion and decision support systems. Research efforts on such dedicated Situation Awareness (SAW) systems have concentrated on assisting the operator in managing the current situations. However, little focus has been so far on integratively supporting the different phases of *knowledge management* in *SAW systems*, which encompasses the *acquisition*, *representation*, *validation*, *maintenance* and *reuse* of knowledge gathered for and during the use of these systems, such as configuring and maintaining suitable situation templates and exploiting already assessed situations. If operators and domain experts are not supported in these tasks, however, this may discourage them from a successful adoption of such systems in real-world control center applications, as user studies revealed. Based on these, and the lessons learned from the application of our SAW system implementations BeAware! and CSI to the domain of road traffic control, we therefore propose a first step towards a tool suite fostering knowledge management in SAW systems, which stretches from the configuration phase of the system to its runtime maintenance in the light of evolving environments and user needs.**

I. Introduction

Situation Awareness Systems. Human operators in today's control centers (e.g., air or road traffic control) need to monitor a plethora of information obtained from diverse sources. To assist them in detecting critical situations within this information flood and taking timely actions, they need suitable fusion of information and support for decision making, as provided by so-called Situation Awareness (SAW) systems. By reporting event combinations that require the operator's focus, i.e., *situations*, these systems reduce the cognitive load on the operator by targeting her attention towards such critical occurrences. *Template-based* SAW systems [1] allow for the specification and detection of specific *situation types* (STs) of interest and have proven their value for detecting a priori known, recurring patterns in a range of different application domains (e.g., maritime surveillance tasks [2] or road traffic monitoring [3]). In this sense, they are a kind of knowledge-based system, as these STs formulate the knowledge about the sought-after real-world situations and must be provided by domain experts. However, their usefulness hinges on the adequate specification of the STs of interest, i.e., the maintenance of that situational knowledge.

Knowledge Management. Consequently, a need for dedicated specification tools has been recognized from the earliest developments on template-based SAW systems on [4]. However, user studies as well as our experiences with our own SAW frameworks indicated that the application to real-world problem domains requires adequate support beyond mere representation of the situational knowledge, i.e., STs. Firstly, tool support for establishing of the *appropriate* situational knowledge base (KB) is needed. This is since: (i) Domain experts need to transform their potentially tacit knowledge to explicit representations of the situations of interest, but may be hampered by the *knowledge acquisition bottleneck*, as user studies on control center operators revealed [5]. Especially the specification of evolving situations, i.e., the behavior of a situation over time, has been identified to be a difficult task. (ii) Current tools pursue a purely top-down approach, thus lack support for determining the "fit" of the specified situational knowledge to real world data (i.e., *validating* the suitability of the represented situational knowledge), as well as *exploiting* the situational knowledge gathered during runtime, such as the encountered situation instances and tracked operator actions and preferences, which could be reused during situation assessment for refining predictions and action recommendations. Secondly, tool support for *maintaining* the situational KB is needed. This is due to the i) ever-evolving nature of both the environment under control, but also ii) changes of the system and iii) its usage by control center operators, therefore a one-shot *establishment of the situational knowledge* turns out to be insufficient, requiring a *continuous adaption and refinement of the situational KB*. However, this is scarcely supported in current SAW systems, as our survey on the evolution support of SAW systems revealed [6]. Overall, tool support for an integrative situational knowledge management for SAW systems involving the human domain expert in the loop is currently lacking.

Contribution. Therefore, in this paper we introduce our concepts on SEM^2Suite, a tool suite supporting the broader management of situational knowledge for SAW systems. SEM^2Suite assists domain experts in (i) specifying STs

by (ii) fostering knowledge acquisition and (iii) validation regarding potential situations of interest, (iv) maintaining the SAW system, i.e., assure its "fit" to evolving environments and usage needs, and information use and reuse, as it allows to (v) explore and (vi) systematically exploit encountered situations and undertaken actions.

Structure of the Paper. In the next section, we attempt at identifying the different tasks regarding the management of situational knowledge encountered in SAW systems, exemplified on our own previous work on a SAW framework for control center applications. In Sec. III, we discuss related work addressing knowledge management (KM) issues of SAW systems, before introducing SEM^2Suite, a tool suite for KM in SAW systems, in Sec. IV. Finally, Sec. V presents an outlook on future work.

II. Knowledge Management Tasks in SAW Systems

In literature, a series of different KM processes comprising partly different tasks can be found (e.g., [7], [8], [9], [10]) — however for the domain of managing the knowledge relevant for template-based SAW systems, up to now no common understanding has been gained so far. Although Jakobson's framework of Situation Management [11] serves as a valuable basis and addresses some KM-related issues of SAW systems (such as highlighting the central role of the *Situation Model*, and characterizing the concepts of *Situation Acquisition* and the *Situation Memory*), its aim is on characterizing the different processes and tasks necessary in SAW systems. However, no further characterization is provided w.r.t. the specific tasks required to build and maintain the situational KB.
Thus, basing on this broad body of literature, in the following we elaborate on those tasks — and their challenges and open issues — we identified as foremost relevant based on findings reported in literature and our experiences on evaluating our own template-based SAW-Systems BeAware![1] [12], [3], [13] and CSI[2] [14], [13] in the domain of road traffic management (RTM), supported by our project partner, Austria's highway agency ASFINAG[3]. In these ontology-driven SAW systems, domain experts transfer their *domain knowledge* to the SAW system by specifying STs representing the sought-after real-world behavior (thus populate the *Model* KB, as shown in Fig. 1), which are then translated to rules. During *situation assessment* (SA), the SAW system's *Inference Engine* matches *data* obtained from the *observed environment* against this rule base. Matched rules trigger the creation of a *situation instance* of the corresponding ST, which is reported to the *control center operator*, and stored in the KB's (Situation) *Memory*.

Knowledge Representation doubtlessly denotes one of the core knowledge-related tasks in template-based SAW systems, which encompasses the modeling of the domain knowledge, usually encoded by means of ontologies, and the specification of the STs of interest in terms of this domain ontology. Thus, this phase requires both a suitable knowledge representation for situations, i.e., situation model for specifying STs, as well as dedicated tools supporting a user-friendly specification of

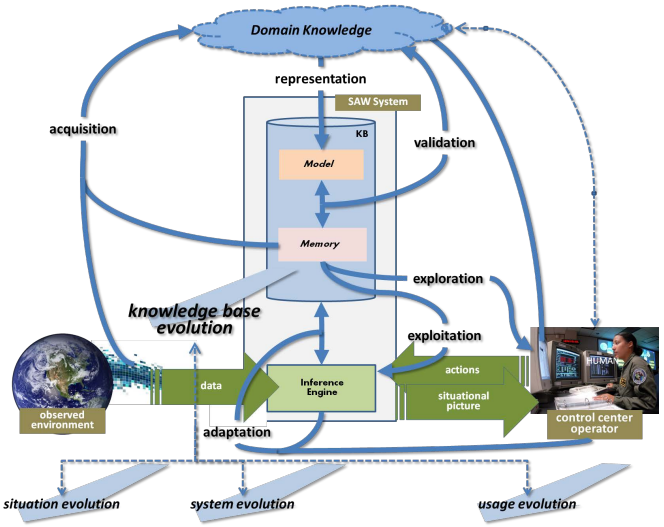

Figure 1: The different tasks for managing situational knowledge in SAW systems.

these STs. Different tools have been proposed which aim at supporting domain experts in specifying STs (e.g., [4], [2], [15]). These mostly employ the JDL situation model [16] as knowledge representation model for STs, which, however, does not provide sufficient account of evolving situations, as has been acknowledged recently [5]. Only [2] supports the specification of *evolving* situations, however, does not allow to characterize the different evolutionary phases (e.g., *trigger*, *climax* and *clearance* phase, as distinguished in [12]), nor criticalities. Whereas highlighting critical, on-going situations supports human operators in mitigating them in a *reactive* way, the ultimate goal of a SAW system should be to provide *proactive* or *prescriptive* support, i.e., alert the operator already if a situation may be *developing towards a critical situation*, thereby, allowing the operator to prevent the situation from escalating. Therefore, BeAware! supported the specification of STs along their different evolutionary phases, i.e., the domain expert could formulate dedicated *trigger*, *climax* and *clearance* STs, which however yielded different situation instances for each phase. During the evaluation of assessed real-world data sets, we learned that in order to enable tracking and analysis of *evolving real-world situation instances*, we would need a *situation evolution model* (SEM) allowing to model situations across their *entire lifecycle*, as proposed in [17], [18].

Knowledge Acquisition. The problem of *Knowledge Acquisition*, i.e., the elicitation of the domain experts' knowledge, has been widely acknowledged in the SAW system community [19], [20]. Prior to being able of specifying STs, the question arises of *what* actually would be STs of relevance in the corresponding application domain. This is complicated by the fact that highly skilled, experienced experts face more difficulties in articulating their know-how than novices, and the more tacit knowledge is, the more valuable it tends to be [21]. This is a crucial factor for template-based SAW systems, which depend on the externalization of domain experts' and operators' know-how. User studies involving control center op-

[1] situation-awareness.net
[2] csi.situation-awareness.net
[3] www.asfinag.at

erators revealed that the problem is even more severe regarding the specification of *evolving situations* [5]. Therefore, it has been concluded that in order to support domain experts in leveraging their tacit knowledge, suitable means for knowledge discovery and data mining facilities should be provided [19], [20]. Although data mining techniques have already proven their capability of discovering important information for SAW applications, existing tools, however, require expertise in data mining. Therefore, domain experts would need to be provided with preconfigured analysis functions [22].

From our experience with BeAware!, we can report that especially the aim to detect evolving situations from their *trigger* on introduces novel challenges: Whereas *climax* situation states are typically known, their preconditions are often not known in sufficient detail, indicating the need for supporting an *incremental acquisition* of the required situational knowledge, starting from the specification and assessment of *climax* situation states and providing a "rewind" functionality allowing to step back in time to investigate on their preconditions.

Furthermore, in our evaluation domain of RTM it turned out to be completely unrealistic to assume that every potential situation of interest can be specified a priori, as we need to expect to encounter novel, previously unknown situations at runtime. Therefore, analogous to the human operator, the SAW system must be capable of reacting to the unexpected. However, whereas (data-driven) anomaly detection based techniques for situation assessment have proven their superiority in these tasks (e.g., [1], [23], [24]), and hybrid SAW systems have been proposed that include both rule-based as well as anomaly detection based situation assessment modules [25], [26], current SAW systems do not provide means to transfer the results obtained with anomaly detection methods to rule-based situation assessors, as suggested for future work in [26]. This would allow to convert the unanticipated to novel knowledge, or refine existing STs (e.g., if previously unknown preconditions or evolutions are detected), thereby maintaining and enhancing the situational KB. Therefore, in the light that recent information fusion models emphasize the need for an integrated approach of combining data analysis with template-driven approaches [27], concepts for supporting the acquisition of novel knowledge based on data analysis techniques are needed.

Knowledge Validation. Whereas *knowledge representation* issues have been addressed by providing various ST specification tools, we have not encountered tool-supported means to *validate* the specified STs. However, our experience in the application of BeAware! to real-world control center settings indicated a need for considering these issues. Domain experts should be enabled of validating the modeled STs w.r.t. their *syntactical* and *semantic* correctness (i.e., the specification tool should provide means to check that the relations specified within a ST are non-contradictory). As our experience with BeAware! showed, however, even a semantically correct ST is of limited use if it fails to capture the real-world behavior of interest, which we term the *adequacy* of a ST. For instance when formulating a ST "An accident causes a traffic jam", our system delivered doubtfully few hits — fewer than would be expected by a domain expert. As a matter of fact, this was due to improper specifications, as the mental model of the domain expert did not fully reflect the behavior observed in the data, requiring to adapt and assess the specifications several times. Therefore, we concluded that domain experts would need to be provided with means to assess whether their mental models "fit" the observed data, i.e., means to *validate* their a priori domain knowledge w.r.t. available domain data sets.

Knowledge Adaption. To retain their usefulness over prolonged application periods while interfacing dynamic environments, SAW systems also need to account for various kinds of evolution: Due to the underlying *environment's evolution*, the STs' specificity may change over time, i.e., once specified STs may fail to capture situations of relevance. Operators thus need to detect this *concept drift* [28] and adapt the STs accordingly. Thus, a one-time configuration of STs would be insufficient, as identified in [27], who state the need for a *configuration management of information models*, otherwise the system's usefulness decreases over time. [29] even concluded that a fully automatic system would need to maintain a state similar to SAW on its own, which resembles the idea of enabling self-organization of adaptive systems in response to changing environments. To remain of high value for the operators, the SAW system needs to maintain the maximum possible degree of "fitness", denoting the suitability between knowledge (i.e., the STs) and the problems confronted (detecting and preventing critical situations in a dynamically changing environment). Furthermore, as template-based SAW systems require substantial configuration efforts, means must be provided to capture as much information as possible in a *non-intrusive* way, just by tracking the operator, i.e., allow the system to learn from the operators, as well as infer the operators' preferences and changing usage (*usage evolution*) and adapt towards those. The specification and maintenance of STs could be simplified if the SAW system could autonomously, or in an operator-guided fashion, learn STs and actions of relevance from the operators. This is especially relevant as valuable tacit knowledge often results in some observable action, when individuals understand and subsequently make use of knowledge [21], thereby the tracking of these observable actions allows to dynamically acquire the experience and tacit knowledge of operators. Summarizing, in order to maintain the KB's "fitness" to the operators' tasks, *knowledge adaption* requires the system to address both self-adaptivity as well as the incorporation of user-driven changes at runtime (as also demanded in [27]).

Knowledge Exploration. Operators learn from previously encountered situations and thus gain more experience, which in turn enhances their future decisions. Thus, a SAW system should support *Investigative Situation Management* [11], i.e., provide the operators with means to explore and learn from the constantly growing *situation memory*, thereby allowing to explore also distinct evolution patterns ("Why did a situation evolve in a specific way?", "What would be the typical evolution?"). Regarding our own experience, our domain experts also indicated the desire to assess the success of undertaken actions, which, in case of large-scale events, is conducted by means of after-action-reviews. Forensic analysis w.r.t. to undertaken actions (e.g., "What went wrong?", "What could have been done better?") thus requires a means to assess the effect of actions onto the evolution of encountered situations, thereby requiring that these actions are tracked and linked to the persisted situation evolutions.

Knowledge Exploitation. Furthermore, an intelligent SAW system should be capable of *Situation Learning* [11], i.e.,

exploit the Situation Memory to enhance its future SA and predictions. Currently, learning from previously encountered situations is rarely addressed in rule-based SAW systems, as opposed to data-driven SA approaches, which by definition base on observed data, and are capable of learning if they are continuously incorporating new data (e.g., as in [30]). Moreover, to enhance the system's decision support capability (i.e., support *resolution* [31]) and aid operators w.r.t. to the question "What should I do?", the system should also be able to retrieve those actions that have yielded the most desirable outcome in the past. Thereby, it could capture and exploit the experience of operators by tracking their actions, similar to the approach of experience retrieval suggested in [32].

Summarizing, Fig. 1 outlines the interplay between the different KM tasks, which highlights the complexity of KM in SAW systems. KM in SAW systems is further complicated by the need for a rapid *KB evolution* due to environment, usage, and consequent system evolution. In the light of these challenges, operators and domain experts would benefit from tools supporting the identified KM tasks and addressing the discussed issues.

III. Related Work

In the present section, we elaborate on tools supporting KM issues in SAW systems.

Knowledge Acquisition. Whereas *knowledge acquisition* has been considered an important problem within SAW systems' research (cf. Sec. II), we did not encounter tools specifically addressing this issue within template-based SAW systems.

Knowledge Representation. Several tools have been proposed that aim at facilitating the specification of STs for domain experts, which, however, provide still limited support regarding the specification of STs for evolving situations.
Matheus et al. introduced RuleVISor, a custom SWRL rule editor, which represents a central configuration component of their Situation Awareness Assistant SAWA [4], and allows to model STs in terms of a domain ontology, which are then translated to the corresponding rules.
Edlund et al. also developed a dedicated ST editor, a configuration component within their ontology-driven, rule-based SAW framework [2]. They especially highlight that it allows for the specification of *evolving* situations, by denoting whether the relations employed within a ST need to hold concurrently or subsequently. However, this only allows for the specification of a single evolution path, as alternatives cannot be specified. Furthermore, as a crucial drawback, their rule engine implementation detects such situations only *after* they completed their overall evolution.
Costa et al. proposed SML, a *Situation Modeling Language*, which allows for a graphical composition of STs, which are then automatically translated to JBoss Drools[4] rules [15]. However, whereas they consider the possibility of referring to *current* and *past situations*, a more fine-grained distinction of situation evolution is not supported. Although their approach allows for specifying behavior over time, only a single course of events can be specified. Therefore, it is not possible to specify a *branching* ST, which might evolve either one way or the other.

Knowledge Validation. None of the discussed approaches that allow for the specification of STs provide specific tool-supported means to directy assess one's definitions on available data during specification, in order to examine the validity and specificity of the proposed STs.

Knowledge Adaption. Furthermore, we did not encounter specific tool-based support regarding the maintenance of the (template-based) SAW system, such as detecting *concept drift*: Only Edlund et al. state the need that the system must be able to incorporate both long-term rules, as well as rules that need to be added temporarily, thus are only valid for a certain time frame [2]. However, they leave the maintenance of the temporal validity to the operators, as they do not foresee any automated support to detect when certain rules become outdated over time, or their specificity changes.
As a first step towards bridging the gap between template-based and anomaly detection SA approaches, Rhodes et al. introduced SeeCoast, an SAW system for coastal surveillance that incorporates both, a rule-based situation assessor, as well as data-driven anomaly detection modules [25]. Thereby, this SAW system supports both the detection of known, recurring situation patterns, as well as reacting to novel situations. However, whereas their anomaly detection SA module continuously adjusts itself to environmental changes, by gradually being re-trained on new data, and even learns from operator feedback, they do not discuss means whether and how the *anomalous situations* could be futher on integrated in the existing rule-base, thereby, do not provide any suggestions to evolve the system's KB. A similar approach of combining rule-based SA with anomaly detection techniques has been described in [26], which have been neither integrated, which is, however, suggested as future work.

Knowledge Exploration. Approaches that allow to assess the situational knowledge the SAW system has gathered over time, or allow to explore what the system has learned, are rarely found. Riveiro et al. proposed a visualization interface that jointly visualizes normal behavioral statistical models (generated from the data) and expert rules (specified by domain experts) within a "scatter plot grid", thereby providing a means to explore "knowledge space" of the system [26]. Their visualization aims at exhibiting how well the experts' knowledge correlates with the models extracted from the data.

Knowledge Exploitation. Currently available template-based SAW systems do not attempt at systematically reusing the situation instances accumulating in the situation memory to refine SA or projection. An interesting approach regarding the exploitation of experience from expert operators has been suggested in [32]: Operators' decisions, i.e., actions, are continuously tracked, along with the current context (employing a dedicated *Action-Observation-Hypothesis* model), and stored in a KB, allowing to retrieve these experiences to aid junior operators in similar contexts.

However, whereas interesting approaches have beep proposed addressing specific KM tasks, we did not encounter approaches that comprehensively support different KM tasks in template-based SAW systems, i.e., stretch across *knowledge acquisition*, *knowledge representation*, *knowledge validation*,

[4]http://drools.jboss.org

knowledge adaption, *knowledge exploration* and *knowledge exploitation*. Therefore, as a first step towards providing a KM tool that integratively supports distinct situational KM tasks, we elaborate on a tool suite fostering dynamic KM in SAW systems.

IV. A Tool Suite for Knowledge Management in SAW Systems

In the present section, we propose SEM^2Suite, a modular (situation evolution modeling and maintenance) tool suite, which supports KM tasks in rule-based SAW systems, thus addresses:

Knowledge Acquisition: comprises components suited to explore and mine existing domain data sets, in order to *acquire* insights about potentially relevant situation types, as well as anomaly detection components that aim at detecting not yet specified situations, which may be transferred to the KB.

Knowledge Representation: comprises a specification pane that allows to graphically compose *Situation Evolution Types* (SETs), based on assembling object types and relation types from a suitable domain ontology.

Knowledge Validation: comprises a set of different validation components, that support to validate the previously modeled SETs w.r.t. their *syntactical* and *semantic* correctness, as well as their *adequacy*, i.e., their "fit" to and specificity on real-world domain data sets.

Knowledge Adaption: comprises components that assure self-adaptivity as well as the incorporation of user-driven changes at the runtime of the system, such as components that detect *concept drift* within the monitored environment that influences the specificity of the modeled SETs, components that allow operators to activate, deactivate or modify specific SETs, performance monitoring components that seek to optimize the situation assessor's performance, and user tracking components that aim at optimizing the UI towards the user's preferences.

Knowledge Exploration: comprises (visual) analytics interfaces that allow operators to explore the accumulated situation memory, i.e., allows to address questions such as "How do situations typically evolve?" or "What have been good actions in specific situations?", thereby also supporting forensic analysis.

Knowledge Exploitation: comprises components that aim at reusing the experience accumulated in the situation memory by refining the prediction on currently observed situations and the recommendation of the most promising actions, based on comparing the current situations to similar situations observed in the past, which are retrieved from the situation memory, thereby, specifically supports real-time decisions.

We realized this range of functionality within a flexible, configurable tool suite (cf. Fig. 2), as not every monitoring task may require each aspect, or may provide all prerequisites. For instance the components for *knowledge acquisition*, as well as *knowledge validation* w.r.t. *adequacy*, depend on the availability of a domain data set. If no data set is available, SETs can only be modeled in a top-down fashion, and be fitted towards the data at runtime.
Finally, it should be emphasized that SEM^2Suite abstracts the modeling concepts from concrete technologies and implementation details, thereby allowing to interchange these technologies. In the following, the different components will be explained in more detail.

A. Knowledge Acquisition

SEM^2Suite's ***Data Analyzer*** component represents a container for interfacing various established data mining tools applicable to data from SAW systems, such as R[5] and Weka-STPM [33], however, provides a dedicated user interface tailored towards SAW data sets and applicable, preconfigured analyses. The employed knowledge discovery tools allow to gain ideas about potential STs, by detecting aberrations from the normal environmental picture, which often correspond to situations of relevance [1], rare and unusual events, and often co-occurring object types. Furthermore, STs may be autonomously suggested by means of association rule mining [34], which therefore retrieves objects in specific relations that may serve as initial ST suggestion that can be refined by the domain expert.
The ***Data Analyzer*** further comprises anomaly detection modules that report unusual courses of events at runtime, which may correspond to novel situations. If the operator decides that a reported anomaly corresponds to a novel SET, the ***Data Analyzer*** extracts the object types and relations between the encompasses objects, and provides the operator with an initial SET suggestion, that can be confirmed or refined by the operator, before it is incorporated into the SET KB.

The ***Hotspot Analyzer*** supports the incremental specification of SETs: Operators may begin with specifying the climax, i.e., hotspot, situation state, and let the system assess these hotspot situation states on the provided domain data set. Based on the found situations, the system may automatically try to infer typical preconditions of these (how were the objects being part of the hotspot related in earlier timesteps?), and may already provide the user with a template including the inferred, potential relations and object types. On the other hand, if automatic SET suggestion fails, the user should be enabled to sift through the provided examples (by providing a "rewind" functionality to assess the situation's history) and use her experience to infer the likely causalities.

B. Knowledge Representation

A suitable knowledge representation lies at the very heart of Situation Management (SM) [11], i.e., a model for specifying templates for real-world situations of interest, and allows the tracking and persisting of these situations in order to enable both investigative as well as predictive Situation Management. To address the *knowledge representation* challenges regarding evolving situations reported in Sec. II and account for the dynamic evolution of real-world situations, we therefore proposed a dedicated Situation Evolution Model (SEM) in our previous work [18], which constitutes a knowledge representation that allows to model, track and reason upon evolving situations. Therefore, situation evolution types (SETs) specified by instantiating the SEM for a specific evolving ST of interest denote the basic *knowledge representation* utilized by SEM^2Suite.

Domain Ontology Configurator. The specification of SETs requires a suitable domain ontology. In order to allow for a

[5] http://www.r-project.org

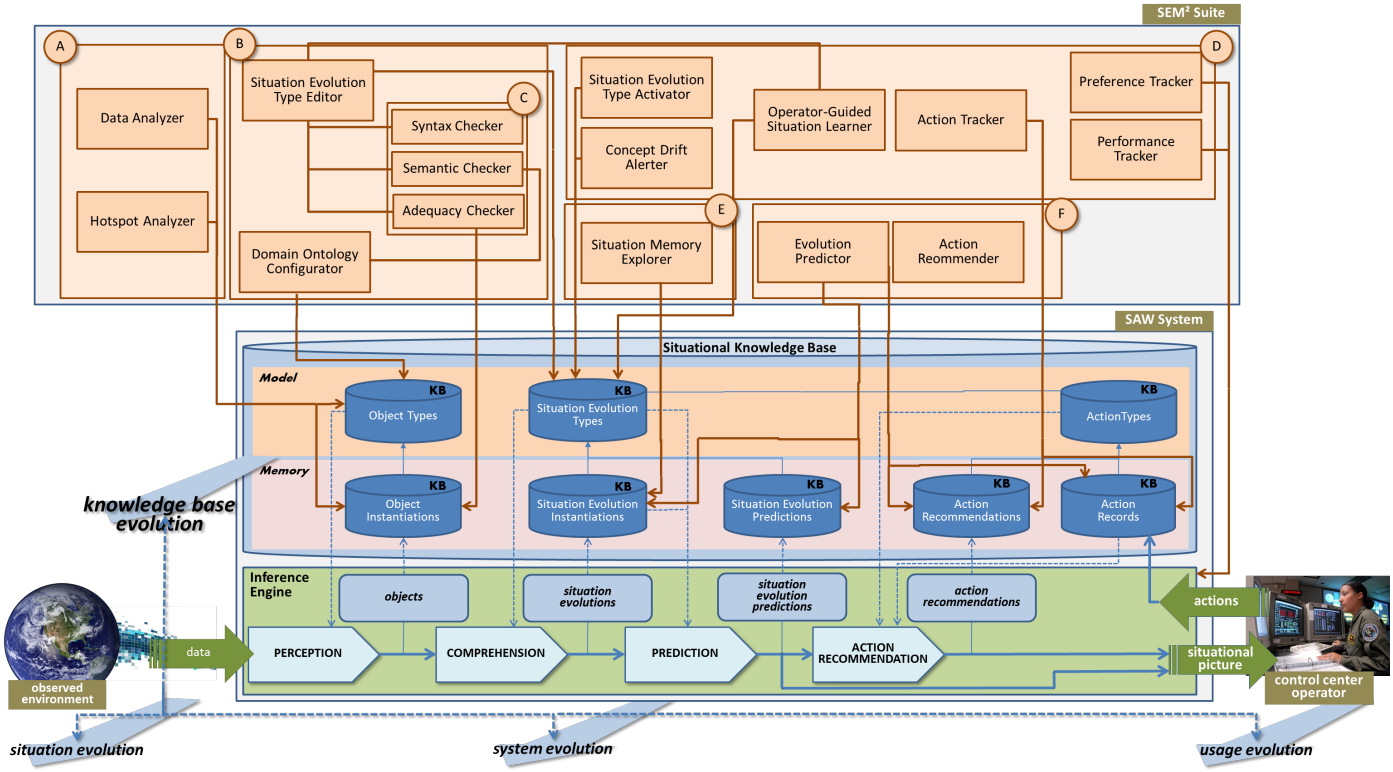

Figure 2: An architectural overview on SEM^2Suite's different components and how these interface and interact with a SAW system's *Situational KB*.

domain-agnostic SA and SEM^2Suite functionality, we therefore require to engineer this domain ontology by extending a generic, i.e., domain independent *SAW core ontology* (based on [35]). This SAW core ontology characterizes the basic entities of relevance in SAW applications, notably *Object Types* sensed from the environment, which are characterized by spatial and temporal properties (e.g., a traffic jam which possesses a specific location and duration), *Relation Types* indicating semantic relationships between such *Objects*, such as spatio-temporal relations (e.g., basing on well-known relation calculi such as Allen's temporal relations [36] or the Region Connection Calculus [37]), and the SEM, which allows to model evolving situations.

Thus, to develop a concrete SAW control center application based upon this framework, the SAW core ontology needs to be extended by a suitable *domain-specific ontology*, in order to allow for sensing and reasoning about the information of relevance of the corresponding domain. Furthermore, the basic spatio-temporal relations provided by the framework need to be configured by *parametrizations* suitable for the given application domain. For instance, suitable threshold intervals for spatial relations such as *Close* and *Far* need to be determined, which may substantially differ between a maritime traffic control center and a highway agency.
In our current prototypical implementation, we follow a UML-based approach of ontology-engineering [38], [39], by employing a dedicated UML modeling tool[6] for generating a Java class library comprising the object ontology, database schemes and corresponding Hibernate[7] mappings. Ontology engineering is thus currently decoupled from SEM^2Suite, which expects to be provided with a Java Archive (jar) comprising the data access layer generated by the employed modeling tool.

SET Editor. The SET Editor allows for the graphical specification of Situation Evolution Types (SETs). Fig. 3 shows a screenshot of our current prototypical implementation of the ***SET Editor***, which features the diagram for the specified SET "Wrong-way driver approaching tunnel" (WDApproachesTunnel). A SET captures the potential evolutions of real-world situations by means of a state-transition system, whereby the states correspond to a particular *relational* state within the evolving situation, for instance the state of affairs where an object of object type "Wrong-way driver" (WD) is in a relation type "Close" to another object type "Tunnel" (depicted by the correspondingly named circles in the diagram shown in Fig. 3). As a monitored real-world situation instance evolves, in subsequent time steps, at one point the relation between the observed objects may change, for instance the monitored wrong-way driver object may come "very close" to the tunnel, which is modeled as a *transition* from the situation state type (SST) "Wrong-way driver close tunnel" to the SST "Wrong-way driver very close to tunnel". In order to reference a specific object across different SSTs, i.e., refer to it throughout its evolution, the different object types must be identifiable by means of an alias (which is unique within a specific SET), as shown in the Outline View in the lower right section of the editor screenshot, which highlights the different object refer-

[6]Visual Paradigm for UML, http://www.visual-paradigm.com

[7]http://hibernate.org

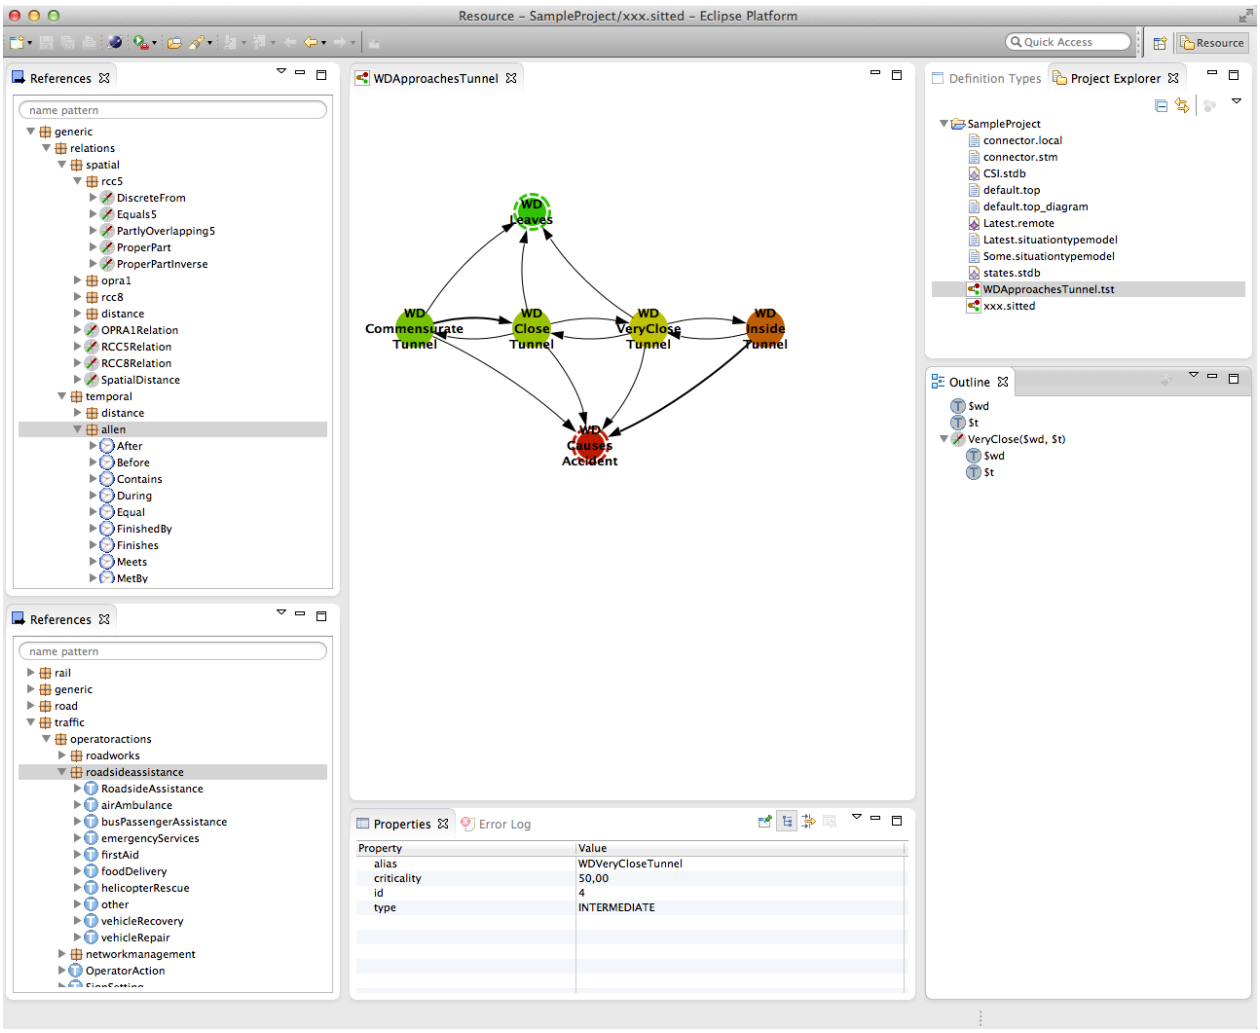

Figure 3: A screenshot of the current SEM^2Suite prototype, which shows the ***SET Editor***.

ences (i.e., object types identified by an alias) and relations employed within a specific SET. A domain experts thus models the different relational states of evolving situations, i.e., SSTs, by assembling them from object types and relation types from the employed domain ontology, which is comprised in the two views on the left side (the upper one comprises distinct relation types, the lower one comprises the object ontology), and specifying the evolution transitions between these SSTs afterwards. Tracked real-world situations then correspond to a specific *path* throughout their SET, i.e., are persisted as a list of interlinked SST instances created by matched SSTs (along with the temporal information how long each of these "situation snapshots" lasted). Furthermore, the domain expert can specify additional properties for those SSTs, such as their evolutionary phase and criticality (visualized by means of the different colors of the states shown in Fig. 3). SEM^2Suite fosters the reusability of modeled components by allowing to include already specified SSTs in other SETs. Furthermore, different action types can be specified for a SST, such as "closing the tunnel for incoming traffic" (and encompassed actions, such as "setting the traffic lights") and "informing emergency units", which can be suggested to the operator as soon as a situation instance of that specific SST is detected.

Translator. Finally, each modeled state needs to be translated to a rule according to the language the SAW system's rule engine employs, such as JBoss Drools Expert[8] or Jess[9]. Therefore, SEM^2Suite requires a dedicated ***Translator*** library that allows to attach the translation rules how a SST pattern can be translated into the target rule language, i.e., a IF-THEN pattern. The left-hand side of the rules (IF) needs to match the Object Types i.e., check the type of the matched object and bind it to the given *alias*. Furthermore, a specific translation for each Relation Type must be implemented.

Optimizer. Since a situation comprises a *set* of relations, no ordering is implied on these relations. Regarding the evaluation of IF-THEN rules, however, the different IF-clauses are processed in sequential order. Thus, the ordering of the IF-clauses matters in the respect that a clause that significantly reduces the number of matched objects should be prioritized, as it raises the performance of the rule evaluation, as we evaluated in [13]. Therefore, a dedicated ***Optimizer*** strategy library

[8]http://www.jboss.org/drools/drools-expert
[9]http://www.jessrules.com

encodes different optimization patterns targeted at specific rule engines and application domains. These may provide rules on the sequential ordering of relation evaluation, such as spatial relations should be prioritized over temporal relations.

C. Knowledge Validation

Assuring *syntactical correctness* means to check that the specification complies with the chosen specification formalism, i.e., the SEM. The ***Syntax Checker*** may validate that Object Types are referenced by an alias, the arity of relations is met, and Object References overlap across a single evolution step, as demanded by the formalism described in [18].

The ***Semantics Checker*** goes one step further by assuring that the specifications correspond to meaningful and possible real-world situations, by taking into account the semantic interpretations of the relations. As it may be possible that specifications are syntactically valid, but semantically wrong, this would not be discovered by a tool providing only syntactical checks. For instance a user may erroneously specify that an accident has occurred *inside* a tunnel and is located *in front of a tunnel*, which is contradictory. However, to enable SEM^2Suite to perform such semantic checks, this information must be provided in the domain ontology, by means of Conceptual Neighborhood Graphs (CNGs) representing the *epistemic knowledge* on how relations of different relation families can be combined (as elaborated on in our previous work on exploiting CNGs for situation projection [40]).

Further validation w.r.t. the *adequacy* of the specified SETs, i.e., determining whether the specified SETs are actually indeed capable of capturing the sought-after real-world situations, can be done by using the ***Adequacy Checker***. This component needs to be provided with a database on data collected from the corresponding application domain, which is accessible from the data access layer loaded in the ***Domain Ontology Configurator***. The ***Adequacy Checker*** can be used to assess *Relation* parametrizations, such as providing histograms for the parametrizations of relation families, and object types of interest. For examples, one may be interested in how the specificity of the SET "Probably fusing traffic jams" defined by the relation "traffic jam j1 *close* to traffic jam j2" changes if the relation type is changes from *close* to *very close*, or if the threshold intervals of these relation types are adapted (e.g., if the threshold interval for the relation type *very close* is changed from 0 - 1 km to 0 - 2 km, thereby widening the set of matched relations). Thus, the ***Adequacy Checker*** pane provides *before - after* histograms, showing the number of situations assessed with the previous and the new definitions. Thereby, it provides a means to inspect and adjust the specificity of the specified SETs, as the domain expert is enabled to incrementally adapt the relation parametrizations and SET specifications, until the desired specificity is reached, depending on whether she wants to receive a multitude of situations, or a smaller set of results. As investigated on in [25], individual operators exhibit different preferences regarding the alarm rate, i.e., situation hit rate. Thus, SEM^2Suite thereby provides operators with a means to configure the SETs' specificity to their personal preferences.

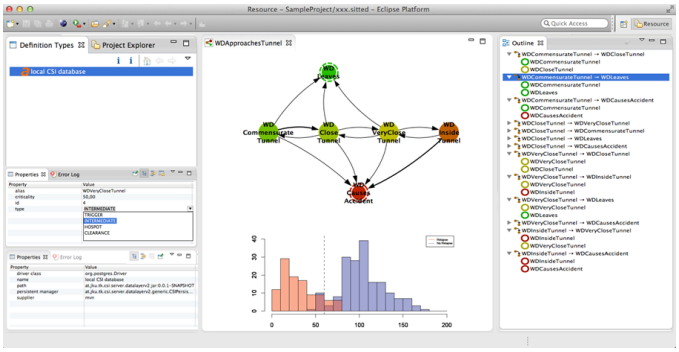

Figure 4: A tentative design draft of the ***Adequacy Checker***, which highlights the quantity structure observed from the data for different definitions (e.g., the quantity structure *before* and *after* changing the thresholds).

D. Knowledge Adaption

Even after a careful specification phase, the specification of the SETs is unlikely to be perfect. Furthermore, as the monitored real-world is continuously changing, these changes need to be adequately incorporated to the SAW system, to retain its usefulness.

The ***SET Activator*** provides an interface for operators to manage the activation and deactivation of SETs at runtime, thereby allowing the operator to enable and disable the assessment of specific SETs. New SETs can be added and existing ones refined, as this component allows to launch the ***SET Editor*** upon a specific SET. However, such adaptions on-the-fly, during runtime of the SAW system, require the SAW system to keep track of the specific version of a SET a situation has been assessed with, therefore, requiring a SET *versioning* system, as well as *provenance tracking*.

The ***Concept Drift Alerter*** aims at preserving the desired specificity of the SETs, in spite of environmental evolution which may induce concept drift. Therefore, it continuously assesses the specificity of the activated SETs, and compares it to their common distributions stored in a KB. Furthermore, this component also stores the distributions of the environmental normalcy models, in order to detect changes within the normal environmental picture indicating concept drift. The ***Concept Drift Alerter*** thus detects if the specificity of a specific SET changes, and reports it to the operator, who is provided with a summary of the encountered changes and thus can investigate on these changes and adapt the SETs accordingly.

The ***Performance Tracker*** aims at providing a continuous runtime tuning of the SAW system, i.e., addresses *system evolution*. By performing a fine-grained tracking of the quantity structure of matched objects and relations, depending on temporal and load contexts, this provides the basis for a detailed analysis on the optimality of the translated rules. For instance, it may have been assumed during configuration of the system that assessing the temporal relations before assessing the spatial relations would be faster than vice versa. However, the logs provided by the optimization tracker may reveal that assessing the spatial relations first would more significantly reduce the number of matches for the next clauses. Thus, the ***Optimizer***'s strategies can be adapted accordingly and existing

SETs may be recompiled thereafter.

SEM^2Suite's ***Operator-guided Situation Learner*** analyzes the human operators' operations, such as the semantic grouping of objects forming a situation for reporting purposes, derives object types and computes relations, and provides the domain expert with these SET suggestions, which can again refine these suggestions.

The ***Action Tracker*** records the actions undertaken by the operators, both routine monitoring actions, as well as actions that represent a dedicated response to a specific situation, which are directly linked to the corresponding situation. Over time, the operator's experience can be thus collected by the system, and allows to asses which actions have proven to be beneficial in which situations.

The ***Preference Tracker*** logs the operator's interaction with the UI, in order to infer the operator's preferences and adjust the OP display accordingly (for instance by only showing the operator's preferred layers and level of detail).

E. Knowledge Exploration

The ***Situation Memory Explorer*** provides *Visual Analytics* support for interactively exploring the *Situation Memory*, thereby allowing to assess the evolution patterns of actually observed situations. The user can define a context of interest by moving a spatio-temporal selection frame and different filters. The ***Situation Memory Explorer*** then constructs the corresponding query and retrieves situations matching the context of interest from the memory. To highlight their evolution patterns, these situations are mapped onto the graphical representation of their SET (i.e., by retrieving the SSTs of their situation snapshots), thereby highlighting the common evolution patterns of the current context of interest, i.e., which developments are more likely and occur more frequent than others, how long situations last on average etc. By modifying the context, for instance by panning and zooming the spatial selection frame (cf. Fig. 5), the user can assess how these evolution patterns may change under different contexts. The ***Situation Memory Explorer*** thus allows to save a snapshot of the current context and result, and allows to arrange different assessed contexts in a scatter-plot like fashion, in order to compare and contrast the results. For instance, it may become apparent that situations in tunnel X tend to develop less critical than situations in tunnel Y and take less time to be resolved. Based on such insights, for instance, the causalities for these distinct evolution patterns may be investigated on, thereby helping to elaborate on countermeasures.

F. Knowledge Exploitation

The ***Evolution Predictor*** refines the projection of the currently observed situation based on retrieving similar situations observed in the past. By mapping these past situations onto the SET, the transition probabilities from the current situation's SSTs to succeeding SSTs can be computed. The ***Evolution Predictor*** thereby allows to employ different *situation similarity measures*, such as retrieving all situations of the corresponding SET, limiting the results to situations with a similar spatio-temporal context (e.g., the same weekday, the same time of day, the same region), or providing a weighting strategy strengthening the influence of recently

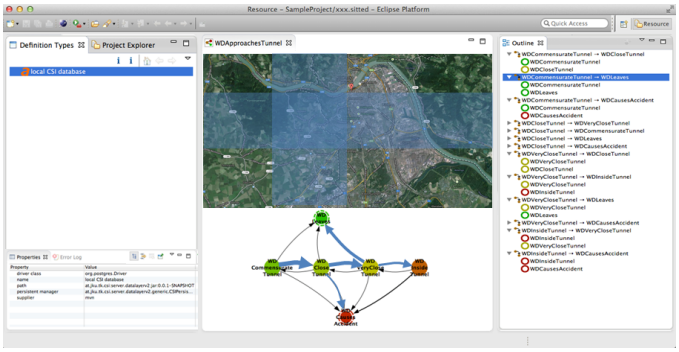

Figure 5: A design sketch of the ***Situation Memory Explorer***, showing the spatio-temporal selection pane which allows to highlight regions and periods of interest. The *Evolution Analysis Pane* (at the bottom of the main pane) highlights the evolution transition frequencies w.r.t. to the selection, as assessed from the *situation memory*, thus allowing the operator to investigate on and contrast local characteristics.

observed situations over situations in the further past (thereby taking environmental evolution into account). Since it may be unclear which similarity measure and projection strategy provides the most accurate results, the ***Evolution Predictor*** thus needs to track its predictions (stored in the *Situation Evolution Prediction KB* shown in Fig. 2) and later on compare it with the actual developments, in order to be able to optimize its prediction strategy.

The ***Action Recommender*** realizes a kind of case-based reasoning functionality, in order to suggest the best action for a current situation. It retrieves situations similar to the current situation from the situation memory, analyzes which of the actions performed on those (retrieved from the *Action Records KB*) have led to the most favorable situation evolution, and recommends these to the operator. The operator is also enabled to inspect these historic situations, if she requires further details on how the past situations evolved. Action recommendations are stored in a dedicated KB (cf. Fig. 2), which allows to optimize action recommendations w.r.t. the operator's preferences (by comparing action recommendations with actually performed actions).

V. Conclusion and Future Work

In the present paper, we characterized the challenges of knowledge management in SAW systems, based on findings in literature and experiences gained from the application of our SAW systems BeAware! and CSI to the domain of road traffic management. Based on the identified requirements, we sketched our vision on a tool suite fostering knowledge management in SAW systems, which especially focuses on incorporating the dynamic aspects of evolving environments and user needs. Regarding future work, we aim at completing our prototypical implementation of this tool suite, and plan to evaluate the feasibility of our approach in a real-world case study involving control center operators from our demonstrators in the domain of traffic management.

VI. Acknowledgments

This work has been funded by the Austrian Federal Ministry of Transport, Innovation and Technology (BMVIT) under grant FFG FIT-IT 829598, FFG BRIDGE 838526, FFG BRIDGE 832160, WTZ/ÖAD AR18-2013 and UA07-2013.

References

[1] R. Laxhammar, "Anomaly detection for sea surveillance," in *FUSION 2008*.

[2] J. Edlund, M. Grönkvist, A. Lingvall, and E. Sviestins, "Rule-based situation assessment for sea surveillance," *Proc. SPIE*, vol. 6242, 2006.

[3] N. Baumgartner, W. Gottesheim, S. Mitsch, W. Retschitzegger, and W. Schwinger, "BeAware!—situation awareness, the ontology-driven way," *International Journal of Data and Knowledge Engineering*, vol. 69, no. 11, pp. 1181–1193, 2010.

[4] C. Matheus, M. Kokar, K. Baclawski, J. Letkowski, C. Call, M. Hinman, J. Salerno, and D. Boulware, "SAWA: An assistant for higher-level fusion and situation awareness," in *Proc. of SPIE Conference on Multisensor, Multisource Information Fusion*, 2005.

[5] M. Nilsson, J. van Laere, T. Ziemke, and J. Edlund, "Extracting rules from expert operators to support situation awareness in maritime surveillance," in *Procs. of FUSION 2008*.

[6] A. Salfinger, W. Retschitzegger, and W. Schwinger, "Maintaining situation awareness over time - a survey on the evolution support of situation awareness systems," in *Procs. of the 2013 Conf. on Technologies and Applications of Artificial Intelligence (TAAI 2013)*. IEEE Computer Society, 2013.

[7] W. R. Bukowitz and R. L. Williams, "The knowledge management fieldbook," *Business Digest*, 1999.

[8] M. W. McElroy, *The New Knowledge Management: Complexity, Learning, and Sustainable Innovation*. KMCI Press, 2003.

[9] M. H. Meyer and M. H. Zack, "The design and development of information products," *Sloan Management Review*, vol. 37, 1996.

[10] K. Wiig, *Knowledge management foundations: thinking about thinking. How people and organizations create, represent, and use knowledge*. Arlington and Texas: Schema Press, 1993.

[11] G. Jakobson, J. Buford, and L. Lewis, "A framework of cognitive situation modeling and recognition," in *MILCOM 2006*.

[12] N. Baumgartner, W. Retschitzegger, W. Schwinger, G. Kotsis, and C. Schwietering, "Of situations and their neighbors—evolution and similarity in ontology-based approaches to situation awareness," in *Proc. of the 6th International and Interdisciplinary Conference on Modeling and Using Context (CONTEXT)*. Springer, 2007, pp. 29–42.

[13] N. Baumgartner, S. Mitsch, A. Müller, W. Retschitzegger, A. Salfinger, and W. Schwinger, "A tour of beAware! a situation-awareness framework for control centers," *Journal of Information Fusion, Elsevier*, vol. accepted for publication, 2014.

[14] ——, "The situation radar: Visualizing collaborative situation awareness in traffic control systems," in *Procs. of the 19th World Congress on Intelligent Transport Systems*, 2012.

[15] P. D. Costa, I. T. Mielke, I. Pereira, and J. Almeida, "A model-driven approach to situations: Situation modeling and rule-based situation detection," in *EDOC 2012*, pp. 154–163.

[16] James Llinas, Christopher Bowman, Galina Rogova, Alan Steinberg, Ed Waltz, and Frank White, "Revisiting the JDL data fusion model II," in *FUSION 2004*, pp. 1218–1230.

[17] M. Kokar, J. J. Letkowski, R. Dionne, and C. Matheus, "Situation tracking: The concept and a scenario," in *MILCOM 2008*.

[18] A. Salfinger, W. Retschitzegger, and W. Schwinger, "Staying aware in an evolving world — specifying and tracking evolving situations," in *2014 IEEE International Inter-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA)*. IEEE, 2014, pp. 172–178.

[19] J. Salerno, M. Hinman, and D. Boulware, "Building a framework for situation awareness," in *FUSION 2004*, vol. I.

[20] S. Das, "Addressing the knowledge acquisition issue for model based level 2/3 fusion," in *FUSION 2007*.

[21] K. Dalkir, *Knowledge Management in Theory and Practice*. Taylor & Francis, 2013.

[22] M. Hadzagic, M.-O. St-Hilaire, S. Webb, and E. Shahbazian, "Maritime traffic data mining using R," in *FUSION 2013*, pp. 2041–2048.

[23] A. Dahlbom and L. Niklasson, "Trajectory clustering for coastal surveillance," in *FUSION 2007*.

[24] B. Ristic, B. La Scala, M. Morelande, and N. Gordon, "Statistical analysis of motion patterns in ais data: Anomaly detection and motion prediction," in *FUSION 2008*.

[25] B. J. Rhodes, N. A. Bomberger, T. M. Freyman, W. Kreamer, L. Kirschner, A. C. L'Italien, W. Mungovan, C. Stauffer, L. Stolzar, A. M. Waxman, and M. Seibert, "Seecoast: persistent surveillance and automated scene understanding for ports and coastal areas," 2007.

[26] M. Riveiro and G. Falkman, "Interactive visualization of normal behavioral models and expert rules for maritime anomaly detection," in *Computer Graphics, Imaging and Visualization, 2009. CGIV '09. Sixth International Conference on*, 2009, pp. 459–466.

[27] E. Blasch, A. Steinberg, S. Das, J. Llinas, C. Chong, O. Kessler, E. Waltz, and F. White, "Revisiting the JDL model for information exploitation," in *FUSION 2013*.

[28] G. Widmer and M. Kubat, "Learning in the presence of concept drift and hidden contexts," *Machine Learning*, vol. 23, no. 1, 1996.

[29] L. Niklasson, M. Riveiro, F. Johansson, A. Dahlbom, G. Falkman, T. Ziemke, C. Brax, T. Kronhamn, M. Smedberg, H. Warston, and P. Gustavsson, "Extending the scope of situation analysis," in *FUSION 2008*.

[30] B. Rhodes, N. Bomberger, M. Seibert, and A. Waxman, "Maritime situation monitoring and awareness using learning mechanisms," in *MILCOM 2005*, pp. 646–652 Vol. 1.

[31] B. Mcguinness and L. Foy, "A subjective measure of SA: The crew awareness rating scale (CARS)," in *Proc. of Human Performance, Situation Awareness and Automation: User-Centered Design for the New Millenium*, 2000.

[32] C. Zhong, D. Samuel, J. Yen, P. Liu, R. Erbacher, S. Hutchinson, R. Etoty, H. Cam, and W. Glodek, "RankAOH: Context-driven similarity-based retrieval of experiences in cyber analysis," in *2014 IEEE Intl. Inter-Disciplinary Conf. on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA)*. IEEE, 2014, pp. 207–213.

[33] V. Bogorny, H. Avancini, de Paula, Bruno Cesar, C. R. Kuplich, and L. O. Alvares, "Weka-STPM: a software architecture and prototype for semantic trajectory data mining and visualization," *Transactions in GIS*, vol. 15, no. 2, pp. 227–248, 2011.

[34] V. Bogorny, B. Kuijpers, and L. O. Alvares, "Reducing uninteresting spatial association rules in geographic databases using background knowledge: a summary of results," *International Journal of Geographical Information Science*, vol. 22, no. 4, pp. 361–386, 2008.

[35] C. Matheus, M. Kokar, and K. Baclawski, "A core ontology for situation awareness," in *FUSION 2003*, vol. 1, pp. 545–552.

[36] J. F. Allen, "Maintaining knowledge about temporal intervals," *Commun. ACM*, vol. 26, pp. 832–843, 1983.

[37] D. A. Randell, Z. Cui, and A. G. Cohn, "A spatial logic based on regions and connection," in *Computer*, B. Nebel, C. Rich, and W. Swartout, Eds. Morgan Kaufmann, 1992, vol. 92, pp. 165–176.

[38] P. Kogut, S. Cranefield, L. Hart, M. Dutra, K. Baclawski, M. Kokar, and J. Smith, "UML for ontology development," *Knowl. Eng. Rev.*, vol. 17, no. 1, pp. 61–64, 2002.

[39] K. Baclawski, M. K. Kokar, P. A. Kogut, L. Hart, J. Smith, J. Letkowski, and P. Emery, "Extending the unified modeling language for ontology development," *Software and Systems Modeling*, vol. 1, no. 2, pp. 142–156, 2002.

[40] N. Baumgartner, W. Gottesheim, S. Mitsch, W. Retschitzegger, and W. Schwinger, "Situation prediction nets — playing the token game for ontology-driven situation awareness," in *Conceptual Modeling – ER 2010*, ser. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2010, vol. 6412, pp. 202–218.