Enabling CoIST Users: D2D at the Network Edge

Dave Braines[*], Alun Preece[†], Geeth de Mel[‡], Tien Pham[§]

[*]Emerging Technology Services, IBM United Kingdom Ltd, Hursley Park, Winchester, UK
[†]School of Computer Science and Informatics, Cardiff University, Cardiff, UK
[‡]IBM TJ Watson Research Center, Yorktown Heights, NY, USA
[§]US Army Research Laboratory, Adelphi MD, USA

Abstract — **Rapid but informed decision-making capabilities at lower echelons are fast becoming a necessity in many coalition operations due to the dynamism associated with such environments. In this paper we investigate technologies to assist CoIST (Company Intelligence Support Team) users operating at the network edge in support of military operations. Through an integration experiment we illustrate the impact of such technologies in rapid decision-making situations. The paper describes the technology integration experiment in the context of a vignette and shows how a natural language conversational interface between human and machine agents in a hybrid team is used. The system can capture local information reporting, infer high value information based on background knowledge, automatically raise intelligence tracking tasks and match, rank and propose appropriate assets to tasks, taking into account contextual factors such as environmental and the distributed network conditions. The approach utilizes ontology-based resource matching capabilities and uses a Controlled Natural Language as a human-friendly – but machine processable – language that is expressive enough to serve as a single common format for both human and machine processing. This capability is designed to operate in a lightweight distributed environment at the edge of the network.**

Keywords — *ontologies, sensor-mission matching, controlled natural language, conversational interaction, data-to-decisions*

I. Introduction

The intuition behind the D2D (data-to-decisions) concept [12] is to discover, extract and process relevant data to provide end users with enhanced situational understanding so that informed decisions can be made about evolving situations in a timely manner. The situational understanding is obtained by interpreting a collection of situation-relevant data and information from disparate data sources captured in a variety of contexts. The need for a D2D capability is critical within coalitions – especially military coalitions – where ad-hoc groups are formed to conduct a collaborative set of mission goals in dynamic environments. To achieve effective decision-making in such environments, the parties involved need to share data and information relevant to the mission objectives and fuse them to obtain a common situational understanding in order to better conduct their respective mission tasks. However, obtaining appropriate situational understanding is a challenging and computationally hard problem. It is exacerbated in coalition operations for a variety of reasons: (1) the required data and information is distributed within the coalition network – i.e., no single party may have all the required data for a particular decision-making problem; (2) experts are by nature qualified to interpret information in a particular domain, thus bringing information from multiple domains together is often challenging; (3) mundane and repetitive tasks place load upon human cognitive capabilities, reducing available cognitive power for harder and more critical problem-solving tasks; and (4) the depth in the communication chain may reduce the effectiveness of the information provided to the end users due to time delays, reduction in specificity, and so forth.

In typical military coalition operations, the end users are the small units (e.g., company, platoon or squad level units) often operating in dynamic constrained environments with limited D2D resources, unreliable networks, and intermittent communications. A key requirement for these small units is the provision of timely and relevant situation awareness data and information about their area of operations whilst also enabling them to maintain a global view of the mission context at a level of detail appropriate to their role. Thus, there need to be mechanisms which support small units in such a way that they can effectively and efficiently gather, process, exploit, and disseminate information about evolving situations and make informed decisions about those situations. This capability should be available without the need to request or rely on data and information from a central authority (e.g. a tactical operation center). Appropriate machine agents to support users in an environment such as this should be able to make intelligent inferences based on the network conditions and knowledge available within the network in a timely manner, and be able to capture some of the experts' reasoning processes. This enables sound and automated (or semi-automated) decision-making to be achieved for some of the more repetitive and low-level aspects of the decision-making process. Any solution to such problems must also address the distributed nature of the information and provide the capacity for configurable decision support agents to maximize system flexibility and ensure that rapid reconfiguration to support emerging threats and new domains of interest are possible.

We have been conducting collaborative research on *Distributed Coalition Information Processing for Decision Making* within the International Technology Alliance (ITA) research program [13]; one key goal of the program is to investigate mechanisms to assist coalition decision makers through automated or semi-automated systems in distributed information environments. In support of this goal we recently demonstrated the applicability of different ITA technology components in an end-to-end D2D technology integration experiment in the form of a demonstration associated with the US Army [2]. This paper presents details of this D2D demonstration and the research propositions on which it is based upon; specifically those aspects intended to enable human-machine interaction in support of D2D capabilities. The demonstration system described here is designed to support intelligence and surveillance operations at the edge of the network with small units of operation such as the CoIST.

The remainder of the paper is structured as follows. Section II defines the research propositions on which the demonstration is based. Section III provides details of the approach, introducing key aspects of previous research used within the solution. Section IV provides a detailed walkthrough of the technical integration experiment that formed the demonstration. A discussion on related and future work is presented in Section V and the paper is concluded in Section VI with some final remarks.

II. Research Propositions

In this section we outline the main research propositions that have been driving our work in this area and have specifically influenced the technology integration experiment. These propositions are motivated by the desire within the D2D context for flexibility and agility, and arise from our previous experiences and ongoing research.

<u>Natural English interaction</u>

The conversation between the human users and machine agents are initiated using natural English language.

This provides a very natural and intuitive environment for the human users, with machine agent responses being returned in Controlled English (CE) [1] which is a formal and unambiguous version of English, enabling the human users to determine with confidence whether the machine interpretation of their request is correct. We believe that this natural English interaction is very important from a usability perspective since Controlled Natural Languages like CE are "easy to read but hard to write" and the overhead to the casual user in writing valid CE sentences would be prohibitive in this setting.

<u>CE is the sole information format</u>

All system components (human and machine) shall use the same information representation format: CE.

Even in cases where the human user states information in natural English it is interpreted by their local machine-agent (their "personal digital assistant"), converted into CE, and optionally presented back to the user for their confirmation. This system-wide usage of CE as the only information representation format creates a powerful homogenous layer that is easily amenable to changes in usage or configuration as the system evolves over time. New models can be added, new machine-agents developed and deployed and new local knowledge sourced from the human agents can be easily integrated into this environment to support evolving and unpredicted needs in addition to more traditional sensor information [14].

<u>Existing data-sources can be used</u>

Existing models and data sources can be consumed in their existing format without any changes required.

CE is an experimental research-grade format with limited tooling and support; it is unlikely that many systems today or in the near future will natively support this specific information format. One of the key capabilities of the language and associated agents is to allow existing data sources and ontologies to be described (in CE), thereby allowed them to be accessed or converted into CE by specialised machine agents that use these descriptions. Whilst this capability is not a core focus of the current research activities we were able to quickly demonstrate the real-time conversion of nearly 4,000 existing concepts from 16 separate OWL ontologies into the CE environment without any modifications to the ontologies

In addition to these key research propositions we also have some important but more applied practical considerations that we used in development of the technology demonstration experiment. These included:

- <u>Human interaction efficiency is key</u>

  We believe that the more "seamless" the human experience is, the more likely they are to use the system. This motivated the investigation into voice-to-text and text-to-voice capabilities on top of the basic text based platform which is the focus of our research; it has enabled us to envisage a more wearable environment in future iterations, for example building on emerging platforms such as Google glass.

- <u>Seamless access to local and global knowledge</u>

  The human user of the system should not need to worry about where the underlying models and data sources are in order to perform their job. The system itself should enable such data sources to be pulled in automatically when relevant and if available from a network perspective. This motivated the separation of concerns for different machine agents, for example the local "personal digital assistant" agent focused on enabling the human user to be able to talk to the system, verses the broader "back office" agents with access to different data sources able to consume incoming data and analyse it in the context of background intelligence and other sensor data as required.

These research propositions and other more practical considerations are revisited again in Section IV in the context of the technical integration experiment.

III. Approach

The technology integration experiment draws together research from three main areas to demonstrate the potential for a human/machine hybrid system to allow local knowledge capture (soft sources) and sensor data integration (hard sources), machine-agent assistance, and dynamic mission tasking. The three research areas are that are brought together to achieve this demonstration are:

- Controlled Natural Language [1] as an information representation format [1]
- Sensor matching for missions [4]
- Conversational interaction between human and machine agents [9]

Out of these three main research areas, two are described in further detail in the following sub-sections:

A. SAM: Sensor Assignment to Missions

Earlier research into asset allocation acknowledges that there are typically multiple ways to achieve a task using sensor-provided data. Using the NIIRS (National Image Interpretability Rating Scales) framework [7] as a basis the SAM (Sensor Assignments to Missions) research characterizes various kinds of ISR tasks that can be achieved using different types of sensing capabilities (visible, radar, infrared, multispectral, acoustic, and so forth). Another key aspect of the SAM research was the development of an approach to automatic task-asset assignment founded on the Military Missions and Means Framework (MMF) [15]. SAM defines models of task and asset types, as well as an automatic procedure for matching tasks to assets. Originally this capability was developed using OWL ontologies, but in our more recent work the representational basis for SAM has been migrated to the Controlled English (CE) language where we refer to ontologies as *conceptual models*.

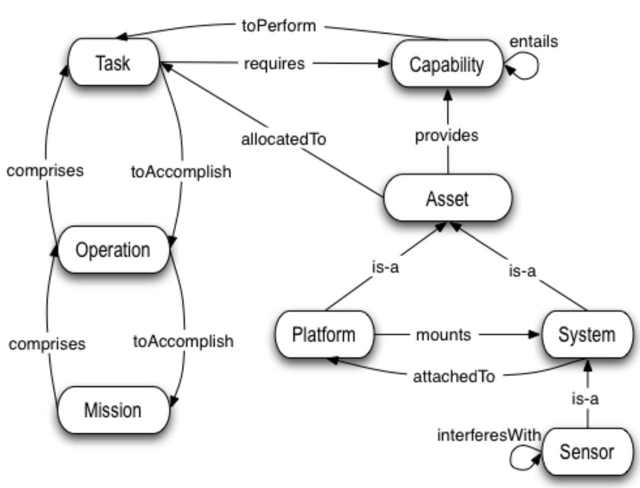

Fig. 1. Missions and Means framework ISR model [4]

Figure 1 shows part of the task-asset matching model that builds on concepts and relationships from the MMF tailored specifically to the ISR (Intelligence, Surveillance and Reconnaissance) domain. In this model missions are comprised of operations that are comprised of tasks. Tasks require capabilities, which are provided by assets. Assets include platforms and systems; systems – including sensors – are mounted on platforms. The relationship *allocatedTo* captures that an asset is assigned to a particular task.

Task types are defined as a pair < *IC, DS* > where *IC* is an intelligence capability (one of: detect, distinguish, identify or localize) and DS is a set of detectable things relevant to the domain of interest (for example kinds of vehicle or building). Examples: *detect {tank}* and *distinguish {tank, jeep}*.
The full details of the reasoning process are described in our previous research [11], however a couple of key features are:

- Some types of task imply others of a lower specificity, e.g. the ability to identify something implies also the ability to detect it.
- The detectable themselves (e.g. buildings/vehicles) are also represented in a hierarchy, supporting the implication that if a given detectable can be detected with a specific capability then all further specialisations of that detectable can also be detected. For example detecting a car means that jeeps, SUVs and pickups can also be detected.

There are simple relationships between sensors, platforms and assets that can be used to create rich and complex combinations known as *bundle types*. Bundle types contain sets of sensors mounted on a specific platform. The sensors and the platform are always compatible and express their aggregate capabilities for use in the reasoning processes.
Bundle types can then be matched to tasks, where tasks are expressed in terms of requirements whilst the capabilities of the bundle types are expressed in terms of sensor capabilities, all expressed within the CE language. The NIIRS framework is the mechanism by which the alignment between tasks and capabilities is calculated. A bundle type is described as minimally matching a given task when no sensor type can be removed from it such that the *matches* relationship still holds[2].
Using this approach SAM can be used to determine all possible matches between user-specified tasks and complex combinations of sensor and platform types in the form of bundle types. Having identified matching bundle types the system will then also identify matching *bundle* instances, for example from a known inventory, taking into account relevant contextual factors such as where, when and for how long the capability is required. To aid the decision-making process a utility function is computed for each matching bundle type. The earlier SAM research also focused on algorithms to efficiently match multiple tasks to multiple bundles in a resource-constrained environment.

[1] Specifically we are using a Controlled Natural Language named "ITA Controlled English" which has been developed as part of our earlier research

[2] Whilst this paper deliberately avoids giving explicit CE examples (please refer to references for extensive examples of these) all of the terms in italics in this section and in figure 1 are CE relationships. For example a valid CE sentence that may define a task using this model is:
there is a task named t1 that requires the intelligence capability detect and is looking for the detectable thing 'wheeled vehicle' and operates in the spatial area 'North Road' and is ranked with the task priority medium.

Figure 2 shows an overview of the SAM approach. The process operates clockwise from the top left of the figure. The system (or user) specifies a task to *localize SUVs* in a particular area. Internally the matching procedure uses the models and knowledge base to determine that assets with a *Visible NIIRS rating of 4* (or higher) or an *Acoustic NIIRS rating of 6* (or higher) can achieve this task. These ratings are provided by: a *UAV* platform with *video camera* sensor, or an *unattended ground system* platform with an *acoustic array* sensor.

Bundle instances are now identified and ranked with a utility function for this specific task. A bundle instance may contain more than one instantiation of the bundle type, where more than one set of deployed assets is needed to achieve the task. Based on the ranking process the results are presented back to the user for their confirmation. This aspect of the process could be automated too, choosing the highest utility solution, or the first above a certain threshold, but for this demonstration we wanted to highlight the power of semi-automatic allocation with the human and machine agents working together as part of a hybrid team.

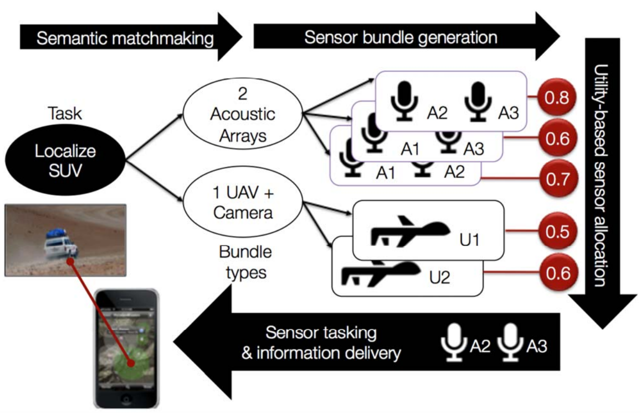

Fig. 2. Overview of the SAM (Sensor Assignment to Missions) approach.

B. Human/Machine Conversations

The second main research area that is used to underpin this technology integration experiment is that of a conversational natural language mechanism between human and machine agents. In this demonstration this takes the form of a human patrol user reporting local knowledge and assisting with the asset allocation decision. For human/machine conversational interaction we focus on supporting three main kinds of interaction using the CE language:

Human->machine interactions where the purpose of the interaction is to mediate between Natural Language (NL) and CE forms of human-provided content to allow an unambiguous but human-friendly conversation to unfold. The human initiates an interaction in NL and the machine feeds back CE, prompting the human to confirm or refine the CE and agree the unambiguous meaning.

Machine->human interactions where the purpose of the interaction is for the machine to initiate the conversation, specifically to inform a human or ask them for information. While it is possible to use CE for this purpose, it is often more convenient to present the "gist" of full CE in a more compact and human friendly form. This may be ambiguous but will often be simpler for the human to read. As a protection against harmful ambiguity or to simply fulfill a desire for precision from the human user they are always able to ask "why" to a gist form of a message in which case the full CE version is returned. In addition to CE content, communications may have other kinds of linked content, for example imagery or a reference to a document.

Machine->machine interactions where the purpose of the interaction is to exchange information between software agents. The conversation is carried out through an exchange of CE content that is directly machine processable. Using CE as a uniform information representation across all agents whether human or machine avoids miscommunication problems and also provides a very straightforward mechanism for human agents to "listen in" or be copied on machine agent conversations when appropriate.

In our earlier research we conceptualise a conversation as a series of *cards* exchanged between agents, including humans and software agents. These cards contain text, which can be in NL or CE. To support human-machine conversation we allow three kinds of card content: fully-natural language, formal CE, and a form of template-based CE that provides the gist of complex sets of CE sentences for brevity and easier human-readability. Through this mechanism a conversation unfolds through a series of primitive *communicative acts*; for example, queries, assertions, or requests [8].

CE cards are broken down into various types, each of which supports a particular communicative act:

- *ask card* that contains a CE query;
- *tell card* that contains CE statements other than queries (e.g. facts);
- *confirm card* that contains CE content derived from the content of a preceding NL card;
- *expand card* that requests the formal CE form of the content of a preceding gist card;
- *why card* that requests an explanation (CE rationale[3]) for the content of a preceding ask or tell card.

A *conversation* is a sequence of cards exchanged between agents, with causal relationships between each pair of consecutive cards. Following [8], we define *conversational policies* as rules that describe permissible conversations between agents, specifying allowed sequences of cards and constraints on the attributes and content of individual cards

[3] In CE "rationale" is the way in which the premises for any inferred conclusions are communicated to the user, and are implemented in the language via the "because" keyword.

depending on their place in a sequence. Figure 3 shows the sequence rules for the card types we have defined.

This conversational capability using the CE cards model is the underpinning for all communication between agents in this technology integration experiment. User interface components use the underlying CE knowledge base for rendering situation awareness information as well as interacting with that information and other users via the CE card conversational mechanism.

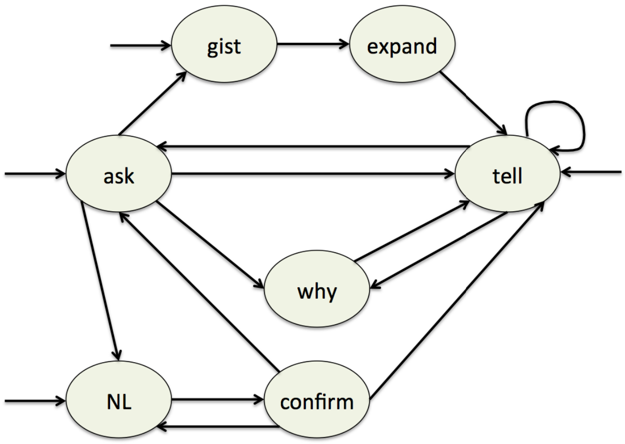

Fig. 3. Conversation sequence rules for CE-Cards [10].

IV. Technology Integration Experiment

The following terms are used in this section:

- **CoIST** – Company Intelligence Support Team
- **SAM** – Sensor Assignment to Missions; as described in detail in Section II.
  There is also a machine agent named "Sam" that provides these capabilities within the demonstration.
- **RTSA** – Real-time situation awareness application, running in a browser in the CoIST environment.
- **MoIRA** – Mobile Intelligence Reporting Application, running in a browser on a mobile device in the field with a patrol user.
  There is also a machine agent named "Moira" that is the local "personal digital assistant" to the patrol user and engages in conversational with the user.

The technology integration experiment described in this paper is based around the first part of a surveillance and asset-tasking vignette defined in our earlier research [10]. The basic goal is the demonstration of semi-automated matching and assignment of appropriate sensor assets. This is based on real-time information coming in from existing sensors and human patrols operating in an area of interest and encountering a potential High-Value Target (HVT) vehicle. The vignette identifies a storyline against which the demo is given, starting with the reporting of a suspicious vehicle by a human patrol users (reported via a mobile device, using spoken voice with voice-to-text conversion capability). This user engages in a brief conversation with the system to confirm the details of their observation.

The system then uses background knowledge (sourced from [5]) to infer that the vehicle sighting reported by the patrol user is a potential HVT since the plate is associated with a known HVT. The system automatically raises a task and uses sensor/mission matching capability to propose suitable sensor choices to the CoIST user via RTSA that is showing the unfolding information in the area of interest.

The vignette is defined at an operational level appropriate to a CoIST team and defines interactions in a the specific domain of patrol reporting, intelligence analysis and high-value target tracking, however the underlying components are flexible and not specific to this vignette. For example, the sensor/mission matching capability can be used to match any sensors to any missions assuming that they are defined with semantic features that support the sensor matching algorithms.

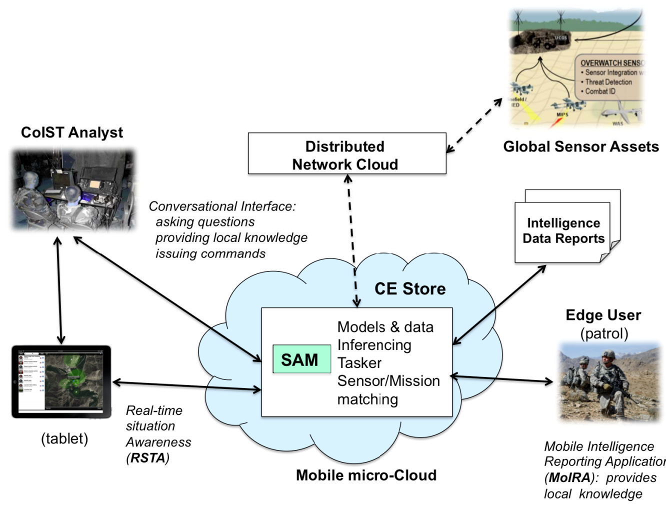

Fig. 4. System context diagram for the demonstration.

The underlying CE language can be used to define any semantic model(s) in any domain(s) of interest in a manner akin to the use of OWL ontologies, with the main difference being the use of human-friendly language [6]. Finally the conversational interaction is based around a set of models and agents also defined in the CE language to enable conversation to occur between agents (human or machine). Thus, this generic capability can be used to allow any conversation on any topic to flow between any set of agents and human users regardless of whether the agents facilitating the conversation "understand" the contents of the conversation. This flexibility on all fronts is key to our research interests and is directly motivated by requirements of a D2D environment.

A. Storyline

Following closely the vignette described earlier, the storyline for the demonstration is:

1. CoIST RTSA environment is providing general background knowledge and situation awareness to the CoIST user in the Forward Operating Base (FOB).
2. Patrol user in the field reports, via NL, their sighting of a suspicious vehicle with various information including license plate. They can speak or type this.

3. Moira (the patrol user's personal digital assistant) processes the message to see whether it can be understood in the context of the current operation and responds with a CE representation showing the interpretation of what the patrol user said.
4. Patrol user confirms (or corrects) this CE interpretation and submits this back to Moira.
5. Moira confirms receipt and passes the new information (in CE) on to Sam, the machine-agent local broker for company level operations.
6. Sam has access to many more data-sources than Moira and fuses data from background intelligence information with the incoming field intelligence to infer that the reported vehicle is related to a high value target that may be operating in the area.
7. Sam reports this by creating (in CE) a new HVT sighting and reporting this to Moira plus any other agents in the vicinity. These agents can then advise their human partner of this new information if it is appropriate to their task or current interests and if policy (optionally also defined in CE) allows.
8. Sam also raises a new task (in CE) to localize and subsequently track the vehicle associated with the HVT sighting and consults available assets in the area to determine those that are best matched to carry out the task. In making this assessment Sam takes into account any number of relevant additional factors that might affect asset performance; in this case weather conditions and available bandwidth.
9. Having compiled the list of matching assets and computed their utility score based on core capabilities with respect to the requirements of task and local factors (weather and bandwidth) these candidate assets are then routed to the CoIST user back at the FOB and presented as an alert on the RTSA application.
10. The CoIST user is then able to use their judgment to pick the best asset for the task[4]. For example they may have better local knowledge about current weather than the weather sensors and may use a visual sensor that has downgraded utility due to fog in the area when the human user knows that the fog has just lifted.
11. Finally, the asset selection is fed back (using CE) to Sam, and a tasking of the physical assets to perform the required task occurs.

The full vignette itself continues with various other interactions in a similar vein, but these are not covered here.

B. Datasets

Various pre-existing models (ontologies) and datasets were used as part of this demonstration. In a real operating environment most of these would likely be provided by existing entities beyond the CoIST in a variety of formats and structures including static sources and live APIs or data-sources. Some might be commercially available (or Open Source) whereas others may be provided by other teams or from higher echelons. In all cases the data and models used in this demonstration are unclassified, often publically available.

The CE language and associated tooling allows users to quickly and easily develop their own models, or extend the models of others; for example in order to capture new types of incoming information that have suddenly become relevant to the current task. We assert that the use of CE as the unified information representation language for all parts of the system and for human and machine agents facilitates such extensions and makes them more accessible to system users who are not IT specialists or knowledge engineers. This potentially enables the system to be more flexible and supportive of unexpected changes as during operations, rather than a traditional stove-piped system designed for a single operational purpose.

The following datasets are relevant to this demonstration and can be represented in CE either by defining them in CE directly, or by describing existing non-CE datasources and transforming these in real-time into CE:

- Sensor catalogue, local (provided by COIST)
- Sensor catalogue, from higher echelon
- Mission details (provided by CoIST)
- Weather information
- Terrain information[5]
- Sensor / mission matching algorithms (SAM)
- Situation awareness[6] information
- Output: sensor-tasking requests (assignments) for either local or higher-echelon resources

In addition to these data sources we envisage that in a real operation the CoIST team would engage with the environment in terms of making decisions, providing relevant local knowledge and guidance to improve the system.

C. Functional Components

The following functional components comprise the technology integration experiment operating environment:

<u>Controlled English Store (CE Store)</u>

This component provides the core CE processing capabilities that underpin all the system components including the machine-agents and their interactions with the human users. Within the CE Store all of the models, data, rules and commands are defined (in CE) to support the required capabilities, including: reporting of local knowledge, asset catalogue, mission definition, matching of assets to missions, as well as domain-relevant background knowledge such as intelligence information. In addition to the main user interface components (described below) and the machine-agents the CE

[4] This is semi-automated assignment in that the machine compiles and ranks the valid options but the human user makes the final decision. This can easily be fully automated with the highest utility match being automatically selected if desired.

[5] Currently only rendered on the mapping layer; not explicitly converted to CE entities

[6] i.e. Information from numerous sources relating to the area of operation, current mission(s) of focus for the COIST, background intelligence etc.

Store also provides a generic "Engineering Panel" interface to allow debug-level interaction with all of the CE components.

Various deployment configurations are possible with the CE Store, including running multiple separate CE Stores within a given server, as well as distributing functionality, models and knowledge across multiple CE Stores or servers. For this demonstration one CE Store is used to support the "Moira" agent who is the personal digital assistant to the patrol user in the field and a separate CE Store is hosted "back at base" (or "in the cloud") in the CoIST infrastructure to support the "Sam" agent who has access to additional data-sources as well as access to the mission and asset catalogues and local environmental conditions via various sensors.

<u>RTSA application</u>

Running in a laptop/tablet environment using browser technology to provide a real-time situation awareness picture using map-based rendering of CE data from numerous sources relevant to the CoIST area of operation. This component is extensible in real-time as new concepts and information are defined in the CE Store; it has been implemented to detect and render additional concepts if they are defined and is not implemented to only handle specific pre-defined concepts. The main mode of operation for the RTSA application is one of passive information presentation to maintain situation awareness for the CoIST user, enabling them to navigate around the underlying information model as appropriate and also add new local knowledge if needed.

In situations where human input is required (e.g. to confirm which asset should be used to fulfil a specific task requirement) the RTSA application alerts the user to the need, providing the context of the request and any supporting information as well as the list of possible choices ranked according to their perceived utility and with the option to look at the rationale (explanation) as to why the asset is being offered and how the utility was computed.

<u>MoIRA application</u>

This is the field-user application and is designed to run on an android mobile phone or tablet device. Since the interaction is based around a conversation and our demonstration includes both voice-to-text and text-to-voice capabilities it is also feasible to run this application on a wearable device (such as Google glass) where the user does not require any form of screen based application in order to interact with the system since they can use their voice and their ears to interact with the Moira agent to convey and confirm their reports. In this demonstration we did not use any wearable technology but this has been demonstrated in earlier research [9].

The MoIRA application supports the reporting of field intelligence from the patrol user via natural language text and confirms the interpretation of the text using unambiguous CE allowing the user to optionally confirm (or correct) the message before propagation to other systems. The application takes advantage of device-level information that is relevant to the report such as the date/time, geo-location and identity of the user making the report. These meta-data properties are also converted into CE and passed across the system.

In the current demonstration the MoIRA application only supports the reporting of unstructured plain text messages, however in previous research [1] we have demonstrated the automated processing of imagery from mobile phones and the creation of more structured SPOT/SALUT[7] format messages.

V. FUTURE WORK

We see the work described in this paper as a first step towards creating a hybrid human-machine environment to support complex collaborative behavior, however there are many additional factors and research directions that can be taken into account to improve the capability further, for example:

Intelligence analysts are increasingly used to working in modern collaboration environments and with social networking tools. Wollocko et al. propose an approach that exploits collaboration using familiar concepts from social media, enabling analysts to identify decision-relevant data distributed across databases and residing in the mental models of colleagues [16]. Reinforcing the idea that including social collaboration can improve the quality of intelligence analysis, [3] offers a web-based application to collate imagery of a particular location from media sources to provide an operator with real-time situational awareness. This work provides evidence that not only it is useful to collaborate within the same analyst team, but also when collaboration is extended to the crowd and mediated by an intelligent software agent, the outcome of the intelligence analysis can be greatly improved.

One key aspect of our work is the ability to use a human-friendly language such as CE to serve as the human and machine basis for the models, knowledge base and inferential aspects of the system. Whilst the promise of this approach is good, one key capability that is a clear requirement for success in this space is the rapid consumption and alignment of models and data sources in other more popular formats. It is not a credible position to assert that potential users must create (or convert) any existing resources into the CE form in order to use a system such as this. Instead we have started some investigations into real-time mapping and conversion of such sources into CE at runtime as required. For the demonstration described here a number of existing OWL ontologies were analysed and converted into the CE format using a simple OWL-to-CE converter agent within the CE Store environment. At this stage this work is not complete however the initial activities have identified that this is a plausible undertaking, building on techniques from the ontology alignment literature. Our initial investigation focused both in terms of semantic expressivity (i.e. can the OWL models be converted to CE) and runtime performance (i.e. can this be done in real time). In this future work we will be drawing further upon existing research into ontology alignment techniques to determine different patterns for this process and will also retain our

[7] See http://armypubs.army.mil/doctrine/DR_pubs/dr_a/pdf/fm6_99.pdf

fundamental focus on the use of CE as a human-friendly language in which the conversion process can be described.

VI. Conclusions

Our ongoing research into CE as a pervasive information representation format for both human and machine users has been coupled with the earlier Sensor Assignments to Missions (SAM) research and demonstrated here against a vignette that shows a hybrid human/machine environment. The demonstration shows the natural language conversational interaction between human and machine agents and demonstrates the reporting of local knowledge (the suspicious vehicle sighting), the inference of high value information from this when fused with background knowledge (the inference of the HVT sighting) and the automatic raising of a task to track the vehicle in question, with the human user selecting the asset to be used from a ranked set of options taking into account the core task requirement, the asset capabilities and relevant environmental factors such as weather at the target location and available bandwidth for sensor data flow once deployed.

All of these capabilities are built onto the CE infrastructure, enabling a pervasive single language to be used for all aspects of the knowledge representation environment, including: models, facts, rules and commands. This is for all aspects of the solution across multiple domains, including: sensors and assets, missions and tasks, detectables, conversational interactions, weather and bandwidth conditions as well as the core domain of interest which in this case is the tracking of high-value targets in an area of interest.

During informal discussions with stakeholders and subject matter experts we have ascertained that the model based asset/task assignment research encapsulated in the "SAM" research is deemed to be of significant interest to the community today. The CE language and conversational interaction between human and machine agents was also of great interest and perceived to be a potentially disruptive technology in this space, however the relative maturity of the language and the lack of an acknowledged associated standard makes this a longer term proposition for likely adoption.

Acknowledgment

This research was sponsored by the US Army Research Laboratory and the UK Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the US Army Research Laboratory, the US Government, the UK Ministry of Defence or the UK Government. The US and UK Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon. The authors would like to acknowledge contribution from the ontology team at the US Army Communications-Electronics RD&E Center (CERDEC) Intelligence & Information Warfare Directorate (I2WD) in the development of the CoIST vignette use case. Also, our thanks go to Chris Gwilliams from Cardiff University and Flavio Bergamaschi from IBM UK for their assistance in the development of the technology integration experiment infrastructure.

References

[1] Braines, D., Mott, D., Laws, S., de Mel, G., & Pham, T. (2013, May). Controlled english to facilitate human/machine analytical processing. In *SPIE Defense, Security, and Sensing* (pp. 875808-875808). International Society for Optics and Photonics.

[2] Braines, D., Preece, A., de Mel, G., Pham, T., (2014, May). Agile sensor tasking for COIST using natural language knowledge representation and reasoning. In SPIE Defense, Security and Sensing (In press). International Society for Optics and Photonics.

[3] Brantingham, R. and Hossain, A., "Crowded: a crowd-sourced perspective of events as they happen,"in *SPIE Defense, Security, and Sensing, 2013, p. 87580D – 87580D*.

[4] Gomez, M., Preece, A., Johnson, M., de Mel, G., Vasconcelos, W., Gibson, C., Bar-Noy, A., Borowiecki, K., La Porta, T., Pizzocaro, D., Rowaihy, H., Pearson, G., and Pham, T., "An Ontology-Centric Approach to Sensor-Mission Assignment," in *Proc 16th Int Conf on Knowledge Engineering and Knowledge Management (EKAW*), 2008, pp. 347–363.

[5] Graham, J. L., Hall, D. L., & Rimland, J. (2011, July). A COIN-inspired synthetic dataset for qualitative evaluation of hard and soft fusion systems. In*Information Fusion (FUSION), 2011 Proceedings of the 14th International Conference on* (pp. 1-8). IEEE.

[6] Huijsen, W. O. (1998, May). Controlled language–an introduction. In Proceedings of CLAW (Vol. 98, pp. 1-15).

[7] Irvine, J.M., "National Imagery Interpretability Rating Scales (NIIRS)," in *Encyclopedia of Optical Engineering*, Marcel Dekker, 2003, pp. 1442–1456.

[8] Labrou, Y., and Finin., T., (1998). Semantics and conversations for an agent communication language. In *M. N. Huhns and M. P. Singh, editors, Readings in agents, pages 235–242*. Morgan Kaufman, 1998.

[9] Pizzocaro, D., Parizas, C., Preece, A., Braines, D., Mott, D., & Bakdash, J. Z. (2013, May). CE-SAM: A conversational interface for ISR mission support. In *SPIE Defense, Security, and Sensing* (pp. 87580I-87580I). International Society for Optics and Photonics.

[10] Preece, A., Braines, D., Pizzocaro, D., & Parizas, C. (2014). Human-machine conversations to support multi-agency missions. *ACM SIGMOBILE Mobile Computing and Communications Review*, *18*(1), 75-84.

[11] Preece, A., Norman, T., de Mel, G., Pham, T., (2013, Jan). Agilely Assigning Sensing Assets to Mission Tasks in a Coalition Context. *IEEE Intelligent Systems* 28 (1), 57-63.

[12] Preece, A., Pizzocaro, D., Braines, D., Mott, D., de Mel, G., & Pham, T. (2012, July). Integrating hard and soft information sources for D2D using controlled natural language. In *Information Fusion (FUSION), 2012 15th International Conference on* (pp. 1330-1337). IEEE.

[13] Preece, A. and Sieck, W. R. (2007). The International Technology Alliance in Network and Information Sciences. *IEEE Intelligent Systems, 22(2):18–19.*

[14] Rein, K., & Schade, U. (2012, March). Battle Management Language as a "Lingua Franca" for situation awareness. In Cognitive Methods in Situation Awareness and Decision Support (CogSIMA), 2012 IEEE International Multi-Disciplinary Conference on (pp. 15-21). IEEE.

[15] Sheehan, J.H., Deitz, P.H., Bray, B.E., Harris, B.A., and Wong, A.B.H., "The Military Missions and Means Framework," in *Proc Interservice/Industry Training and Simulation and Education Conference*, 2003, pp. 655–663.

[16] Wollocko, A.B., Farry, M. P., and Stark, R. F. , "Supporting tactical intelligence using collaborative environments and social networking," in *SPIE Defense, Security, and Sensing, 2013, p. 87580E-87580E.*