Agile Sensor Tasking for CoIST using Natural Language Knowledge Representation and Reasoning

Dave Braines[a], Geeth de Mel[b], Chris Gwilliams[c], Christos Parizas[c], Diego Pizzocaro[c], Flavio Bergamaschi[a], Alun Preece[c]

[a] Emerging Technology Services, IBM United Kingdom Ltd, Hursley Park, Winchester, UK
[b] IBM TJ Watson Research Center, Yorktown Heights, NY, USA
[c] School of Computer Science and Informatics, Cardiff University, Cardiff, UK

ABSTRACT

We describe a system architecture aimed at supporting Intelligence, Surveillance, and Reconnaissance (ISR) activities in a Company Intelligence Support Team (CoIST) using natural language-based knowledge representation and reasoning, and semantic matching of mission tasks to ISR assets. We illustrate an application of the architecture using a High Value Target (HVT) surveillance scenario which demonstrates semi-automated matching and assignment of appropriate ISR assets based on information coming in from existing sensors and human patrols operating in an area of interest and encountering a potential HVT vehicle. We highlight a number of key components of the system but focus mainly on the human/machine conversational interaction involving soldiers on the field providing input in natural language via spoken voice to a mobile device, which is then processed to machine-processable Controlled Natural Language (CNL) and confirmed with the soldier. The system also supports CoIST analysts obtaining real-time situation awareness on the unfolding events through fused CNL information via tools available at the Command and Control (C2). The system demonstrates various modes of operation including: automatic task assignment following inference of new high-importance information, as well as semi-automatic processing, providing the CoIST analyst with situation awareness information relevant to the area of operation.

Keywords: data-to-decisions, sensor-tasking, controlled natural language, ontologies, conversational interaction

1. INTRODUCTION

The work described in this paper has arisen from our ongoing research into *Distributed Coalition Information Processing for Decision Making* within the International Technology Alliance (ITA) research program [25]; one key goal of the ITA program is to investigate mechanisms to support the Data-to-Decisions (D2D) concept, specifically to assist coalition decision makers in distributed information environments through automated or semi-automated fusion processes.

D2D is focused on the discovery, extraction and processing of task-relevant data and information to provide end users with enhanced situational awareness and understanding [24]. Agility is a key concern for D2D: Enabling informed decisions to be made about evolving situations in a timely manner. A D2D capability is critical within coalitions – especially military coalitions – where ad-hoc groups are rapidly formed to conduct a collaborative set of mission goals in highly dynamic environments.

Obtaining appropriate situational understanding is a challenging and computationally hard problem, especially in a coalition context, for a number 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 addition to this, end users are often operating at the edge of the network in dynamic constrained environments with limited resources, and unreliable and intermittent network communications

Users in situations such as these could take advantage of machine agents if those agents are agile and flexible enough to operate in such an environment, and be capable of dealing with the issues that would render more traditional knowledge management systems unusable. Such agents should be able to capture some of the experts' reasoning processes to enable 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, removing this burden from the human user. The research focus areas and solution outline described in this paper attempt to address these kinds of issues in this operational context.

The remainder of the paper is structured as follows: Section 2 defines the research focus areas that are drawn together to create the solution infrastructure described in this paper, including details of an earlier vignette that is used to demonstrate the solution. Section 3 gives details of the solution infrastructure itself, with a focus on each of the main components, some Controlled English examples and a brief overview of existing ontologies and other datasets that were used in the demonstration. Planned future work and some comments on related research are presented in Section 4 and the paper is concluded in Section 5.

2. RESEARCH FOCUS AREAS

The system described in this paper draws together a number of distinct ITA research areas. A brief outline of each of these is given below to provide the reader with some context to understand what capabilities each brings to the overall solution. The final sub-section outlines the vignette that we are using to provide use cases for our research.

2.1 Controlled Natural Language / Controlled English

ITA research since 2010 has focused on the use of Controlled Natural Language as a single information representation format to serve as a common language between human and machine agents in a problem-solving context. The specific implementation of this Controlled Natural Language is named "ITA Controlled English" (hereafter referred to as CE) and is based on earlier work by John Sowa that defined Common Logic Controlled English (CLCE) [27]. The research goal for CE beyond previous efforts is to create a pervasive ecosystem in which CE can be used during design and runtime, and in situations involving human-machine, machine-human, machine-machine and human-human interactions. As a language CE is located within a wider spectrum of Controlled Natural Languages [14] ranging from style guides at the informal end, to full structured languages at the formal end.

CE is a subset of the English language with a restricted set of grammar rules and a vocabulary that can be defined and extended by users of the language. CE itself is based on a formal syntax and semantics [16], and is intended to provide both easy readability for human consumption and unambiguous representation of information for machine processing, without the need for transformation to alternative formats. Specifically it provides:

- A user-friendly language in a form of English, enabling assertion of new knowledge, querying of existing information, construction of logical inference rules and execution of specific tasks and processes. The target users include domain specialists (e.g. military planners and analysts) who will often lack expertise in more formal languages [17] [18].
- Precision and formality, enabling clear and unambiguous representation of concepts and their relationships, enabling the construction of rich domain models.
- The CE language embedded as the core component of a wider CE-based ecosystem that enables rules, triggers, simple scheduling and agent invocation. The CE language itself is used as the intuitive means of configuring the processing and reasoning of this wider CE-based ecosystem. This is described further in Section 3.

There are a number of examples of the CE language used throughout this paper, in Section 3, but there is no attempt to define the explicit CE syntax within this paper since the authors believe that the CE material is readily understandable by casual users, at least in terms of conveying the processing outlined here.

2.2 Sensor-Mission Matching

Knowledge representation and reasoning techniques to support sensor-mission matching and assignment have been another focus of the ITA research in recent years, and a thread of research referred to as "SAM" (Sensor Assignment to Missions) has demonstrated the ability to take high-level specifications of information requirements and match these to low level sensor capabilities. Key components of this research utilized existing patterns and ontologies such as the Military Missions and Means (MMF) [26] framework and NIIRS (National Image Interpretability Rating Scale) [13] in order to achieve this key alignment. This is an important capability to support human users of a complex and fast-moving information environment typical in a D2D context. SAM supports a more direct interaction with assets and resources by non-technical users, providing the potential for a more agile and fine-grained allocation process. The original SAM work was built using traditional knowledge management infrastructure including OWL ontologies and reasoners [8]. In more recent research this ontology-based approach has been successfully migrated across to use the CE language with no loss of expressive power or capability [21].

2.3 Human/Machine Conversation

Building on the earlier CE research we realized that having a common human/machine information representation format is a powerful tool, but the human user requires mechanisms through which they can interact with the underlying information in order to make sense of it in the context of their task and to then act upon it or augment it with further valuable "local knowledge". Our initial efforts were in terms of typical information visualisation and exploration capabilities, all of which were based on the underlying CE knowledge-base, but which presented information to the user in a form appropriate to their task, and the type of information they wanted to see [2]. This was a useful exercise in terms of presenting information to the user but it didn't provide a straightforward mechanism to allow the user to explore the information in unpredictable ways, or assert new local knowledge. When considering how humans approach this issue we subsequently proposed a conversational approach to these interactions, primarily to support human/machine interactions, but also, importantly, to support machine-machine and human-human communications using exactly the same information format (CE). This research defines a CE model based on speech-act theory [1] and software agent communications [5] to support the flow of a conversation through the exchange of "cards" containing information in a particular context within the conversation. A secondary aim of this work was also to provide a full "Natural English" interface to the human user, enabling them to carry out their part of the conversation in Natural English, with confirmation of interpretation being returned by machine agents in CE. This approach is predicated on the view that CE is easy-to-read but harder-to-write without training or tooling. When considering the D2D context of this research and the typical users being engaged in existing tasks at the edge of the network we also extended our demonstration of this research to include voice-to-text and text-to-voice capabilities, enabling the use of hands-free components such as Google Glass [23].

2.4 Vignette

A surveillance and asset-tasking vignette defined in our earlier research [24] provides useful context against which the solution infrastructure described in this paper can be understood. The various CE examples that are provided in Section 3 are all based around this vignette and a full description of a technology integration experiment using this same vignette is also available [3]. The vignette provides a storyline against which semi-automated matching and assignment of appropriate sensor assets is shown, utilizing 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 starts with the verbal reporting of a suspicious vehicle by a human patrol user, engaging in a brief conversation with the system to confirm the interpretation of their observation.

Background knowledge (sourced from [9]) is used to infer that the reported vehicle sighting is a potential HVT since the license plate is associated with a known HVT. A task is automatically raised by the system with sensor/mission matching capability (provided by SAM) proposing appropriate sensor choices to the CoIST user back at base.

This vignette is specifically defined for a CoIST team and defines interactions specific to patrol reporting, intelligence analysis and High-Value Target tracking, however the underlying components are flexible and can be used in a wider range of situations. For example: The sensor/mission matching capability and conversational interface can be used to match any sensors to any missions and report information about any domain of interest, assuming that the domain is modeled within the CE environment and any defined sensors and missions are defined with semantic features that support the sensor matching algorithms.

3. SOLUTION INFRASTRUCTURE

The solution that we have built draws together the research focus areas and the vignette described in Section 2 in the context of a number of research propositions and design principles (see [3] for a full discussion) that underpin our ideas about supporting agile D2D in hybrid human/machine teams:

- Natural English interaction
  This research proposition states that the conversation between the human users and machine agents shall be initiated using natural English language.
- CE is the sole information format
  This research proposition states that all system components (human and machine) shall use the same information representation format: CE.
- Existing data-sources can be used
  This research proposition states that existing models and data sources shall be consumed in their existing format without any changes required.
- Human interaction efficiency is key
  This design principle is based on the assertion that the more seamless the human experience, the more likely they are to use the system.
- Seamless access to local and global knowledge
  This design principle is based on the assertion that 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.

We have also taken inspiration from a number of more abstract principles in the development of our experimental infrastructure and the specific user interface components to support the human/machine interactions. For example the CE Store component is implemented as a fairly pure blackboard architecture [12] and for the human/machine interaction components we take inspiration from the Zero Overhead Principle[1] that is prevalent in many mobile apps and is becoming a basic expectation for the users of the "app generation".

3.1 System Components

The system is comprised from a number of different components as well as human and machine agents. The diagram below shows the overall system context and each of the components are described in the remaining sub-sections:

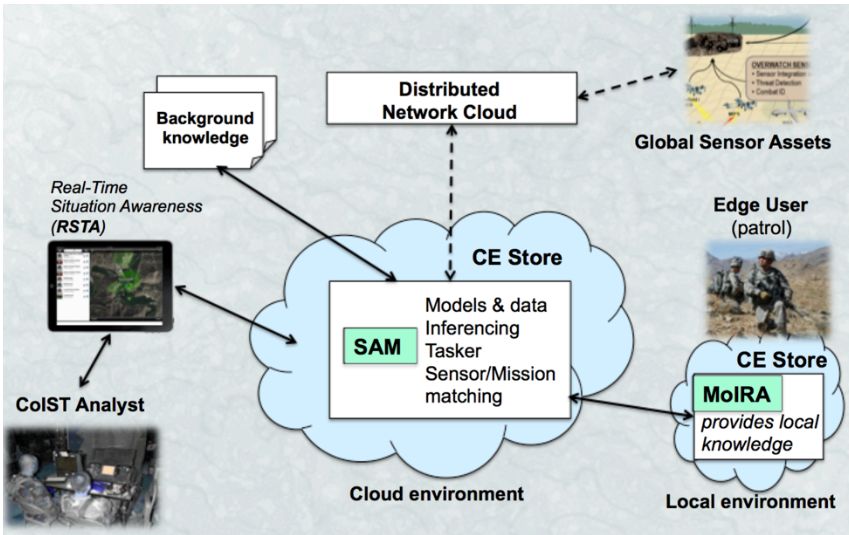

Figure 2: System context diagram for the D2D environment described in this paper

[1] Building For The Enterprise – The Zero Overhead Principle, Dr. DJ Patil – http://techcrunch.com/2012/10/05/building-for-the-enterprise-the-zero-overhead-principle-2/ – checked 5th April 2014.

The RTSA (Real-Time Situation Awareness) application and the CoIST user at the Forward Operating Base (FOB) are not described in detail in this paper but they are included on the diagram to demonstrate the ease with which other users and applications can be integrated into the CE Store knowledge-base environment. This part of the system is described in more detail in [3].

Controlled English (CE) Store

The CE Store[2] is a research-grade implementation of a knowledge-base infrastructure that uses the CE language to provide an agile and extensible environment in which users can develop and extend models of interest as well as information that relates to those models. Queries and rules can also be defined and executed, and the Java-based implementation of the CE Store provides a rich agent API to enable new custom agents to be developed for specific purposes. The agents themselves are defined and configured using the CE language and can be easily shared amongst members of a community. The CE Store can be run in a stand-alone single installation or can be configured to provide multiple CE Stores running within one or more CE Server. The default implementation is as a web application and a rich set of HTTP REST APIs are provided to facilitate integration with existing web-based applications.

Below in Section 3.2 we provide a number of working examples of different types of expression within the CE language.

Triggers and Agents

The CE Store has the ability for anyone to extend the core capabilities through the development of new agents that can access the knowledge-base via a rich Java-based API as described in the previous sub-section. These agents can be instantiated within the CE store, with each being configured for specific tasks or processes as required. For the vignette described in this paper there are two main agents:

- Moira (Mobile Intelligent Reporting Application)
  This agent operates as the local "Personal Digital Assistant" for a human user. Facilitates the interaction with the human user via natural language and confirms the interpretation of any knowledge extracted during the conversation and communicates this new knowledge into the wider environment by passing it to the Sam agent.
- Sam (Sensor-Assignment to Missions)
  This agent handles the fusion of new knowledge from human users (passed to Sam from Moira) into the background knowledge of the system. Sam also raises mission tasks and proposes appropriate solutions for available assets to fulfill these tasks.

These agents and other relevant rules are invoked using a trigger mechanism within the CE Store. Each agent can be registered as a trigger against various events within the CE Store, for example the creation of new instances, or the addition of new information such as properties or relations. In this vignette the Moira agent is invoked whenever a new conversation event is detected (e.g. another human or machine agent engages in a conversation), and the Sam agent is invoked whenever certain information is reported by human users. The various rules within the system (e.g. to check for High Value Target sightings) are defined within the CE Store and executed each time new knowledge is added which may be relevant to these rules. The conversation between the human UK Patrol user and their Moira personal digital assistant, along with the subsequent machine-machine conversation between Moira and Sam is shown in the figure below:

[2] Available for download from the IBM developerWorks site - http://ibm.co/RDIa53 – checked 5th April 2014.

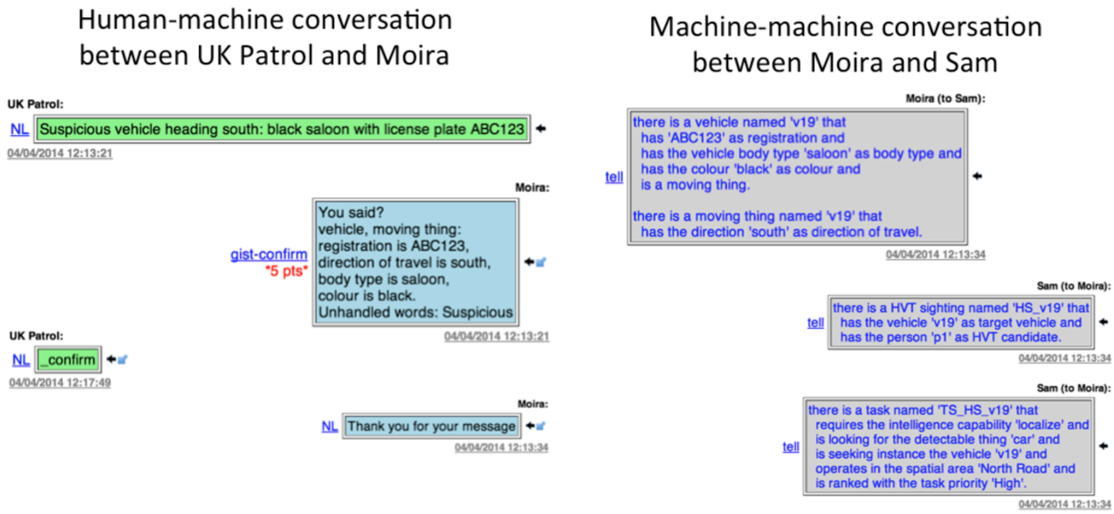

Figure 1: Example of human-machine and machine-machine conversational interaction

As previously discussed, there is a clear relation between these agents and rules and the blackboard architecture. In fact both agents are simply observing the CE Store (the blackboard), watching for information against which they can act. This basic capability is easily extensible with new agents and new concepts as the situation evolves, enabling new specialized agents to be easily plugged into the CE-based system when needed.

Human interaction

The conversational interaction is the main form of human-machine interaction covered in this paper, and the examples given here are all textual. In the demonstration environment we have been experimenting with voice-to-text libraries to allow the human users to report their knowledge verbally, and text-to-voice libraries to enable the Moira agent to vocalize the response back to the user. Whilst this is not a core focus of our research we do believe that it will have a significant impact in any user experiments that we undertake and we have investigated a number of different APIs to achieve good accuracy in interpretation and realistic audio for the vocalization.

3.2 Controlled English Examples

In this section we provide a narrative of relevant parts of the underlying knowledge-base that support the various interactions with the user and between system components. All of these examples take the form of CE sentences that are then described with a brief narrative to outline how each part of the system makes use of the CE information in the knowledge-base to support the conversation. The entire experimental infrastructure is created in the CE Store component and is implemented using the CE language or agents that are invoked as a result of CE activities. All CE in this section is valid operational CE from the experimental environment described in this paper.

Meta-model extensions

Within the CE language we have the concept of a meta-model that can be expressed simply as "*additional facts about the concepts and relationships in the conceptual model(s)*". The meta-model provides a rich and extensible mechanism for providing additional information about the model(s) without needing to extend the core CE syntax every time we wish to add a new feature. In the context of the conversational interaction and Natural Language Processing (NLP) in general we wish to add a couple of key meta-model properties to concepts and relationships to allow us to record synonyms and plural forms. These meta-model extensions are shown in the CE below:

conceptualise the concept C ~ is expressed by ~ the value V.

conceptualise an ~ entity concept[3] ~ E that has the value PF as ~ plural form ~.

[3] "concept" is the generic concept that accounts for all domain concepts, relationships and properties defined in CE, whereas "entity concept" is the specialisation that covers only domain concepts.

These two CE statements add new properties named "is expressed by" (to store synonyms) and "plural form" to the appropriate concepts in the meta-model hierarchy. Having asserted these new properties we are now able to define concepts that are expressed by synonyms and have plural forms (see the next sub-section for examples). This is an extension to the normal CE modeling process where domain models are constructed using concepts and relationships, but without the need for synonyms and plural forms. The addition of this extra information better "binds" the domain models to the words that people use when describing them in natural language and forms a key information source for the conversational agent that is invoked to interpret the natural language sentences coming from the human user during the conversational discussion.

We can also define CE queries to return information from a specific subset of the knowledge-base. CE query syntax is similar to CE fact syntax but wrapped in clauses and using variable names for parameter substitution in a manner similar to technical query languages such as SQL. In future CE language development we hope to support higher-level queries that are more natural to human users but which can be mapped (without ambiguity) to core CE queries such as these.

```
[synonyms]
for which CON and SYN is it true that
        (the concept CON is expressed by the value SYN).
```

The CE query shown above will return all of the synonyms (communicated using CE via the "is expressed by" property as described earlier) for the currently loaded concepts and relationships in the domain model(s).

Creating a domain model

CE examples for the sensor-mission assignment models and the core conversational card models can be found in [21] and [23] and are not repeated here. In this paper we provide some examples of the concepts, synonyms and plural forms created during the main domain modeling process:

```
conceptualise a ~ colored thing ~ CT that has the color C as ~ color ~.

conceptualise a ~ vehicle ~ V that
        is a colored thing and
        has the vehicle body type T as ~ body type ~ and
        has the value R as ~ registration ~.

conceptualise a ~ moving thing ~ MT that has the direction D as ~ direction of travel ~.
```

In the above example we have defined three concepts in our domain model: vehicle, colored thing and moving thing. The intention behind creating the more abstract "colored thing" and "moving thing" concepts is two-fold: (i) it allows us in the modeling exercise to simply state that various concepts (vehicles, clothes, boxes, doors etc.) are colored things, without needing to repeatedly state all the color-relevant properties on each separate concept, and; (ii) it allows the conversational agent at runtime to assert that something which is mentioned with a direction of travel is a moving thing. The ability to have both of these modeling capabilities is important: in the case of colored things all vehicle instances are always colored things (even if they do not state a color) since this relationships is made in the model, against the vehicle concept, whereas a vehicle is not a moving thing unless the instance in question is explicitly asserted to be both a vehicle and a moving thing. This kind of flexibility provides a rich but simple set of modeling capabilities to the domain model creator, in this case allowing them to reflect the "common sense" view of the world that vehicles always have a color, but not all vehicles are always moving.

Having created concepts and relationships in the domain model we are now able to define the synonyms and plural forms for these using the meta-model extensions described previously:

```
the entity concept 'vehicle'
        is expressed by 'car' and
        is expressed by 'truck' and
        is expressed by 'jeep'.
```

```
the relation concept 'moving thing:direction of travel:direction'
        is expressed by 'driving' and
        is expressed by 'heading' and
        is expressed by 'went' and
        is expressed by 'drove' and
        is expressed by 'moving'.

the relation concept 'vehicle:registration:value'
        is expressed by 'license plate' and
        is expressed by 'license' and
        is expressed by 'plate' and
        is expressed by 'registration' and
        is expressed by 'registration number'.
```

There is no limit to the number of concepts, relationships or synonyms and plural forms that can be represented in the CE domain models. In very rich models there will be homonyms (where one word can mean multiple concepts) and in this situation both concepts should state the duplicated synonym via the "is expressed by" relationship, enabling the conversation agent to detect the conceptual ambiguity and attempt to choose the correct interpretation based on the context of the discussion, or present the two interpretations of the homonym back to the human user for their resolution.

In our current implementation we create our domain models and lexical augmentation manually, however in previous work we have established mechanisms for using rich existing ontologies and lexical resources such as WordNet, VerbNet and FrameNet [22] as accelerators for semi-automatically creating appropriate synonym lists in a tool we call the "analyst's helper".

Summarization

A simple mechanism to optionally provide a more human-friendly response than pure CE is referred to as the "gist" form, and this provides a mechanism for customizable summarization by machine agents during the conversation. In the current implementation this is achieved via a very simple mechanism, whereby particular domain concepts are marked as being summarizable things. If they are so marked then any instances of that concept inferred during a conversation will be communicated back to the human user in a gist form, whereas all instances of any other concepts will be returned as plain CE.

```
conceptualise a ~ vehicle ~ V that is a summarizable thing.
        --[4]  (or) ...is a custom summarizable thing.
        --  (or)  ...is a graphical summarizable thing.
```

The example above shows the three supported forms of summarization: standard, custom and graphical. Standard summarization is of the form shown in the vignette conversation where a simple restatement of detected concepts, properties and values is used, whereas custom summarization allows the domain modeler to provide (via CE) a simple template to be used for the summarization activity, and the graphical summarization is an experimental form that is used to allow substitution of imagery to communicate the gist of the message when appropriate, for example when the user may have limited ability to read or listen to long sentences [23]. If a concept is not stated as being a summarizable thing then any instances of it will be returned as plain CE in the conversational interactions.

Background information

In the context of the CE environment "background information" is simply existing knowledge in the knowledge-base (CE Store) and may come from a number of sources and serve many purposes. As previously mentioned we have implemented the CE Store around the idea of a blackboard architecture so background information can be viewed as "writing on the blackboard" and can be on any topic, provided by any agent or user at any time. The example CE below

[4] Note that the "--" marker in CE is the same as a comment marker in some programming languages. This indicates that any text on a line starting with this marker is ignored by the CE parser and in this example it shows the alternative ways of specifying each of the types of summarizable thing.

gives examples of background information that is relevant to the lexical processing of the NL sentences in the conversational interaction, for example common instances such as colors or directions. Although this is background knowledge and is comprised of "CE facts" it can be logically grouped with the domain model since it represents the lexical augmentation of the domain model to support subsequent natural language processing.

```
-- Colors
there is a color named red.
there is a color named black.
there is a color named green.

-- Directions
there is a direction named 'north'.
there is a direction named 'south'.
there is a direction named 'east'.
there is a direction named 'west'.

-- Vehicle body types
there is a vehicle body type named 'saloon'.
there is a vehicle body type named 'sedan'.
there is a vehicle body type named 'SUV'.
```

Another form of background information relates to the various agents that interact in the system, for example the human users and the "Sam" and "Moira" agents. These are also expressed in CE, using concepts and instances creating using the same principles as all other CE information. This means that they are easily extensible with new concepts, properties and information as the situation unfolds, should that be needed. Examples of the CE for these instances is shown below:

```
the service 'Sam' has the country 'UK' as affiliation.
the service 'Moira' has the country 'UK' as affiliation.

the person 'UK Patrol'
        has the country 'UK' as affiliation and
        is located in the spatial area 'North Road'.

the person 'US Patrol'
        has the country 'US' as affiliation and
        is located in the spatial area 'North Road'.
```

The affiliation information captured for these instances enables a very simplistic policy integration point within our demonstration environment, showing that agents with the same nationality will communicate with each other, whereas agents with a non-matching nationality will accept incoming information but not engage in an active conversation (to prevent information leakage). Clearly this is a trivially simple policy rule, but it is simply used to indicate the point in the architecture where a much richer CE-based (or existing external) policy system could be integrated [11] [20].

Finally, the more "traditional" background information takes the form of information directly relevant to the domain model of interest, for example people, vehicles and social relationships. This may be sourced from existing databases, other data sources or APIs, and the role of the CE environment is to provide a consistent human-friendly information format to represent all this information. Also, importantly, this kind of background information can be easily sourced from human users, either by them stating it directly in CE, or by asserting it via a natural language conversation using the conversational interface. This ability to easily and flexibly capture "local knowledge" is a key requirement of a D2D solution, and we believe that the use of CE as a human and machine friendly language helps to support this goal. Some examples of this kind of background information relevant to the vignette are shown below:

```
the person p1
        has 'ABC123' as linked vehicle registration and
        has 'John Smith' as typical name and
        has 'Smithy' as alternative name and
        has the nationality 'British' as nationality and
        has 'shady character' as description and
```

```
        is connected to the person p4 and
        is a high value target.

the person p2
        has 'Bill Jones' as typical name and
        has the cell phone 'x7777 123456' as cell phone.

there is a weather condition named 'windy' that
        has '0.5' as intensity and
        adversely affects the sensing capability 'acoustic sensing'.
```

The CE Store implements a user-extendable trigger mechanism as described previously. The CE example below shows a simple fact statement that defines a triggered event. In this case it is a notification event and ensures that the "Moira" agent (service) will inform the "Sam" agent (service) for every new vehicle instance that is created. This is especially valuable in a distributed environment where the Sam agent may be running against a separate CE Store to the Moira agent, and there may be many Moira agents, any one of which may notify the Sam agent about new vehicles (or anything else).

```
-- Moira will notify Sam of any new vehicles
there is a conversation notify triggered event named 'trig_notify_sam' that
        has 'vehicle' as concept name and
        has 'service' as from concept and
        has 'Moira' as from instance and
        has 'service' as to concept and
        has 'Sam' as to instance.
```

Whilst this notification agent is quite simple it a powerful abstraction when combined with the extensible nature of the underlying CE conceptual model. For example we could change this to notify on any instances of any number of concepts, or create a new abstract concept named "notifiable thing" and then directly assert in CE that specific instances are notifiable things, thereby being much more selective as to what Moira will notify Sam about. All of this is achieved simply by changing the CE example shown above. This very simple notification mechanism can be easily extended to support a richer group-based notification mechanism where agents (or groups) can notify other agents (or groups) based on some more complex query-based matching algorithm. In our vignette this would be especially useful for example to notify all agents in a given area with the correct nationality that a HVT sighting has been made, rather than needing to notify each agent individually.

Finally we have an example of a CE rule that is similar in construction to a CE query except that instead of returning a result set a rule has a conclusion (following the work "then" in the example) that is used to create new knowledge from the rule premises. The example given below is the rule that is used to infer a High Value Target sighting if a vehicle is observed that is linked to a person who is a High Value Target. Any number of rules such as these can be written to take advantage of the semantic features in the model, with both the model and the rules being extensible as the situation develops, as described previously.

```
[HVT sighting]
if
        (the vehicle V has the value R as registration) and
        (the person P has the value R as linked vehicle registration) and
        (the person P is a high value target)
then
        (there is a HVT sighting named HS_#HS_$V[5]) and
        (the HVT sighting HS has the vehicle V as target vehicle) and
        (the HVT sighting HS has the person P as HVT candidate).
```

[5] The exact syntax of the variable for asserting new instances is not described here, but for this example it is suffice to say that this will use the id of the associated vehicle to ensure that only one HVT sighting is raised per vehicle observation.

The result of executing a rule such as the one above is the generation of new knowledge in the knowledge-base, along with "rationale" which is an explanation of the premises that were involved in drawing the conclusion. An example of such a sentence with rationale (following the "because" in the example) is given below:

```
there is a HVT sighting named 'HS_v19' that
        has the vehicle 'v19' as target vehicle and
        has the person 'p1' as HVT candidate
because
        the vehicle 'v19' has 'ABC123' as registration and
        the person 'p1' has 'ABC123' as linked vehicle registration and
        the person 'p1' is a high value target
[HVT sighting].
```

The CE examples given in this section are by no means exhaustive but are included to enable the reader to get a sense of how simple the basic CE syntax is. Using this simple syntax a rich set of models and information spanning multiple domains can be quickly brought together to support human and machine agent teams in an environment where rapid development or evolution of models and situation understanding may be required as the situation unfolds. The agility of the environment also enables the use of initially simple models that can be developed into richer semantic structures as the situation warrants. It also provides a powerful environment for integration existing models and data sources and a brief description of this process is described in the next sub-section.

3.3 Incorporating existing Ontologies

One key aspect to our on-going research and experimentation with the CE language is the human-friendly language serving as the single human and machine information representation format 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 *existing, more widely used, formats*. It is unreasonable to assume that any potential users must create (or convert) any such existing resources into the CE format in order for them to be able 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 work described in this paper a number of existing OWL ontologies were analysed and converted into the CE format using a simple OWL-to-CE converter agent operating within the CE Store environment. Although these are initial experiments the early stage activities we performed have identified that this is a plausible undertaking both in terms of semantic expressivity (i.e. can the OWL models be converted to CE) and in terms of runtime performance (i.e. can this be done in real time). In future work we will be drawing upon existing research into ontology alignment techniques to determine different patterns for this process and will also retain our fundamental focus on the use of CE as a human-friendly language in which the conversion process can be described. The initial results of this conversion exercise for a variety of existing OWL ontologies can be seen in the table below.

OWL to Controlled English Conversion Statistics		
Ontology Name	***#CE concepts***	***#CE properties***[6]
Agent	993	7
AIRS Emotion	78	
AIRS Mid Level	541	
Artifact	302	
BFO	40	
Counter-Terrorism	49	

[6] Note that many of the ontologies do not contain properties since they are used as sources of inter-related concepts modeled with simple parent/child relationships.

OWL to Controlled English Conversion Statistics		
Ontology Name	***#CE concepts***	***#CE properties***[6]
Event	416	
Geospatial	301	
Information Entity	88	23
Information Technology	360	6
OBO In OWL	6	
Quality	684	
RO BFO Bridge	45	12
RO	6	
Time	19	22

Figure 3: A summary of converted CE concepts and properties from various source OWL ontologies

4. FUTURE WORK

The main planned focus for our ongoing research is the investigation of additional capabilities for the conversational interaction between human and machine agents, and we intend to specifically focus on the following capabilities:

- Interjection from machine agents
  To start or redirect the conversation, seeking new knowledge, confirming assumptions, resolving inconsistencies etc.
- Support additional types of interaction
  Question asking, model extension as part of the discussion, multi-sentence narratives etc.

In addition to these specific focus areas within our research we also seek to align aspects of our underpinning work with wider research in the field into at least the following areas:

- Ontology Alignment
  An ontology is intended to represent the shared understanding of a domain [10], however, multiple ontologies could exist for a given domain due to multiple, especially in more agile environments where ontology development and extension by end users is encouraged. The CE language can be used to manually construct ontology alignments, typically through the writing of logical inference rules to propagate information from one ontology to another, however this is time-consuming and prone to error or partial completion. Automated approaches vary from simple syntax-based mapping strategies [19] to those that take advantage of external resources to determine likely matches [6] [15], and more semantically driven approaches [26] [7] that attempt to analyze the structure of the ontologies to determine linkages. We believe that mapping new CE models to a small set of existing higher level CE models will be a task that is understandable by model creators and will provide a degree of ontology alignment capability at the more semantic end of the spectrum, however we intend to review the latest research in the context of our human-led CE-based approach to ontology (model) development.
- Explanations
  Building a richer capability for summarization and explanation of information, inferences or related material is another focus for the conversational interface. In our research so far we have experimented with the simple template-driven "gist" form to provide a more human-consumable version of textual information and have demonstrated some potential for graphical forms of gist as well. The CE "rationale" capability also has a role to play here, however the field of research for explanation and summarization is rich and relevant to our research. According to Buchanan and Shortliffe [4] explanations serve two main purposes for human users of knowledge-based systems:
  (1) they improve the understanding of the knowledge captured by the system among users. Thus, allowing users to efficiently use, debug, maintain, and introduce new knowledge; and

(*2*) they increase the acceptance of the system and will persuade one to try it out. This is because, if the conclusions of a system are shown to be reasonable, users tend to accept the system more.

These areas and those more minor points mentioned in passing through the paper will be considered in the context of our ongoing research, with the aim to demonstrate a compelling human/machine conversational capability that is desirable and efficient for human users, is flexible and agile in re-purposing to new domains and integrates well into existing data sources and ontologies.

5. CONCLUSIONS

We have given details of the research and solution outline for a system to support D2D aspirations for coalition users operating at the edge of the network. Our solution is built upon the human-friendly but machine-processable CE language and demonstrates the potential for a natural language conversational interaction between human and machine agents on top of this underlying architecture. One of the agents (Sam) embodies the earlier ITA research into sensor-mission matching, which has also been demonstrated running in a CE-based infrastructure. An ISR based vignette has been used to demonstrate this natural language conversational interaction between human and machine agents, specifically for the following activities: (i) reporting of local knowledge by human users (the vehicle sighting), including a brief conversation to confirm the interpretation of the reported information; (ii) the automatic inference of new high value information from this field knowledge when fused with background knowledge (the inference of the HVT sighting); (iii) the automatic raising of a task to track the vehicle in question, with another human user selecting the asset to be used from a ranked set of options, and; (iv) the ranking of those 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.

The CE Store provides a pervasive operating environment that is used for the entire knowledge-base, including: models, facts, rules and commands. The CE language is used 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.

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 also 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.

REFERENCES

[1] Austin, J., & Urmson, J. (1975). How to Do Things With Words. Harvard University Press, 1975.

[2] 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.

[3] Braines, D., Preece, A., de Mel, G., Pham, T., (2014, July). Enabling CoIST Users: D2D at the Network Edge. In *Information Fusion* (FUSION), *2014 Proceedings of the 17th International Conference on* (In press).

[4] Buchanan, B.G. & Shortliffe, E.H. (1984). Explanations as a Topic of AI Research. In Buchanan, B. G. and Shortliffe, E. H., editors, *Rule-based expert systems: the MYCIN experiments of the Stanford Heuristic Programming Project*, page 331. Addison-Wesley.

[5] Burg, B., & FIPA, V. P. (2002). Foundation for Intelligent Physical Agents. Official FIPA presentation, Lausanne, February.

[6] Chalupsky, H. (2000). Ontomorph: A Translation System for Symbolic Knowledge. In Cohn, A. G., Giunchiglia, F., and Selman, B., editors, *Proceedings of the 7th International Conference on Principles of Knowledge Representation and Reasoning* (KR2000). Morgan Kaufmann.

[7] Euzenat, J., Meilicke, C., Stuckenschmidt, H., Shvaiko, P., and Trojahn, C. (2009). Ontology Alignment Evaluation Initiative: Six Years of Experience. In *Journal on Data Semantics XV*, volume 6720, pages 158–192. Springer.

[8] 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.

[9] 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.

[10] Gruber, T. R. (1994). Toward Principles for the Design of Ontologies Used for Knowledge Sharing. *Journal of Human Computer Studies*, 43(5/6): 907–928.

[11] Harries, D., Braines, D., & Gibson, C. (2012). Towards an expression of Policy in Controlled English. In the Proceedings of the 6th Annual Conference of the International Technology Alliance.

[12] Hayes-Roth, B. (1983). The blackboard architecture: A general framework for problem solving?. Heuristic Programming Project, Computer Science Department, Stanford University.

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

[14] Kuhn, T. (2010). Controlled English for Knowledge Representation (Doctoral dissertation, PhD thesis, Faculty of Economics, Business Administration and Information Technology of the University of Zurich).

[15] McGuinness, D. L., Fikes, R., Rice, J., and Wilder, S. (2000). The Chimaera Ontology Environment. In *Proceedings of the AAAI-00*, pages 1123–1124.

[16] Mott, D., "Summary of Controlled English," May 2010, https://www.usukita.org/node/1411

[17] Mott, D., Giammanco, C., Braines, D., Dorneich, M., and Patel, D., (2010). Hybrid Rationale and Controlled Natural Language for Shared Understanding. In *Proceedings of the Fourth Annual Conference of the International Technology Alliance*, London, UK.

[18] Mott, D and Hendler, J. (2009). Layered Controlled Natural Languages, In Proceedings of the Third Annual Conference of the International Technology Alliance, Maryland, USA.

[19] Noy, N. F. and Musen, M. A. (2000). PROMPT: Algorithm and Tool for Automated Ontology Merging and Alignment. In *Proceedings of the 17th National Conference on Artificial Intelligence* (AAAI-00).

[20] Parizas, C., Pizzocaro, D., Preece, A., Zerfos, P., (2013) "Managing ISR sharing policies at the network edge using Controlled English", in Proc. *Ground/Air Multi-sensor Interoperability, Integration, and Networking for Persistent ISR IV*, Baltimore, USA, 2013

[21] 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.

[22] Poteet, S., Xue, P., Kao, A., Mott, D., Braines, D., & Giammanco, C. (2013). Controlled English for Effective Communication during Coalition Operations. In *18th ICCRTS "C2 in Underdeveloped, Degraded and Denied Operational Environments"*, Alexandria VA, USA, June 2013

[23] 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.

[24] 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.

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

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

[27] Sowa, J., "Common Logic Controlled English," March 2007, http://www.jfsowa.com/clce/clce07.htm

[28] Zhdanova, A. V. and Shvaiko, P. (2006). Community-Driven Ontology Matching. In *The Semantic Web: Research and Applications*, pages 34–49. Springer.