

The Application of Advanced Knowledge Technologies for Emergency Response

Stephen Potter¹, Yannis Kalfoglou², Harith Alani², Michelle Bachler³,
Simon Buckingham Shum³, Rodrigo Carvalho⁴, Ajay Chakravarthy⁴, Stuart Chalmers⁵,
Sam Chapman⁴, Bo Hu², Alun Preece⁵, Nigel Shadbolt², Austin Tate¹ and Mischa Tuffield²

¹University of Edinburgh, UK, ²University of Southampton, UK, ³The Open University, UK,

⁴University of Sheffield, UK, ⁵University of Aberdeen, UK

Corresponding authors: {s.potter@ed.ac.uk, y.kalfoglou@ecs.soton.ac.uk}

ABSTRACT

Making sense of the current state of an emergency and of the response to it is vital if appropriate decisions are to be made. This task involves the acquisition, interpretation and management of information. In this paper we present an integrated system that applies recent ideas and technologies from the fields of Artificial Intelligence and semantic web research to support sense- and decision-making at the tactical response level, and demonstrate it with reference to a hypothetical large-scale emergency scenario. We offer no end-user evaluation of this system; rather, we intend that it should serve as a visionary demonstration of the potential of these technologies for emergency response.

Keywords

Sense-Making, Decision Support, Semantic Technologies, Emergency Response, Intelligent Messaging.

INTRODUCTION

Modern ‘best practice’ approaches to effective management of the emergency response – for example (HMFSI, 2002) – stress the key role of information as an asset to be managed in this process. With the growth of the internet, the potential for rapidly accessing large amounts of information in emergency response situations is greater than ever before; however, this information is often dubious or untrusted, and so needs to be filtered effectively. The semantic web (Berners-Lee *et al.*, 2001) and associated technologies could provide one solution, in part, to this problem. Through the use of ontologies and semantic annotation techniques, and tools able to reason with these, some of this information processing can be performed automatically. In other words, computers can, to some extent, ‘understand’ the content of information on the web, and so determine its relevance in particular contexts.

In this paper we describe an *e-Response*¹ system that has been constructed to demonstrate how semantic web and other knowledge technologies can be harnessed to assist the response to a simulated large-scale emergency. This system loosely integrates a number of quite disparate technologies; wherever possible, this integration has been achieved at the ‘knowledge level’ through shared ontologies and knowledge bases. From the perspective of the system user – assumed to be a dedicated operator supporting a strategic/tactical response team – these technologies are accessed through an existing ‘sense-making’ tool. Since the system represents an agent in a wider response organisation, decisions taken in the sense-making space can be relayed to other agents, and they, in turn, can communicate issues and information. In this paper we describe the constituents and the use of this system; first, however, we present the scenario used to motivate the development of the system, and its later demonstration.

An Emergency Scenario

The background to the scenario is as follows: “The time is 10am on a Tuesday morning. One hour ago a cargo aeroplane crashed on the City of London [the central business district of London, roughly a square mile in area]. Debris and fuel is scattered over the entire area, and multiple fires have broken out on the ground.” Hence, this is a major emergency, occurring in a densely populated area of a major conurbation, and involving a number of separate incidents that are distributed geographically. Emergency services on the scene have already declared it to be a major incident, and, accordingly, the members of the *Joint Emergency Services Control Centre* (JESCC) have been

¹ *e-Response*: the creation and operation of task-centric virtual organisations to respond to emergency situations.

convened at some location safe from immediate danger². Physically, the JESCC is usually housed within the command vehicles of the emergency services, which also provide access to the necessary communications infrastructure. In addition to senior representatives of the police, fire brigade and ambulance services, typically those personnel assuming ‘global’ (i.e., over the available resources considered as a whole) tactical command of the incident, plus their safety, scientific and press advisors, the JESCC may also call upon representatives of utilities and (local or national) government, and, if circumstances dictate, specialists from industry or academia. We assume that the JESCC is in constant communication with the command centres of the emergency services, namely, the *Fire Brigade Command Support Centre (CSC)*, the *Police Central Communications Complex-Information Room (CCC-IR)* and the *Central Ambulance Control*, each of which has ‘local’ tactical control of its service resources, and that the JESCC is supported by a ‘JESCC Liaison’ team, which handles ‘informal’ communications from the public and media. For the purposes of this scenario we assume the JESCC has the following functional roles:

- to command and control the coordinated tactical response of the three principal emergency services;
- to relay information between the services and act as a conduit for requests for cross-service aid where this cannot be done effectively on the incident ground;
- to respond to incoming information relayed (by the Liaison team) from members of the public and the media.

In this case, the JESCC is supplemented by an additional member: a JESCC computer support officer. It is the responsibility of this officer to perform the following roles in support of the wider aims of the JESCC:

- to make sense of the state of the emergency and the response to it;
- to act as the primary point of contact for information and requests raised from the emergency service command centres, and to relay information and directions from the JESCC to the operational units;
- to provide decision-making support to the JESCC.

The support officer will use the e-Response system to discharge these duties; it is assumed that this officer is familiar with both the system and the location of the emergency. From the perspective of the support officer, the scenario plays out through a series of messages sent to the e-Response system, with each message corresponding to some event occurring during the emergency. For each message, the task of the officer is, simply put, to respond appropriately, that is, so as to move the response organisation towards the achievement of its aims.

E-RESPONSE SYSTEM: TOOLS AND TECHNOLOGIES

In this section we describe the different component tools and technologies (most of which were originally developed in generic forms) individually, before describing how they are integrated to constitute the e-Response system.

3Store and Ontologies

The e-Response system is underpinned by several centralised *3Store* (Harris, 2005) knowledge bases (generically termed *triplestores*, since their contents are represented conceptually as RDF triples) holding shared knowledge. There are a number of OWL ontologies used in the system to represent knowledge in the triplestores (and elsewhere), and hence, used to construct SPARQL queries (through an HTTP interface) to retrieve knowledge from the triplestores. These ontologies define concepts related to the emergency services, buildings, etc., and were generated using material from a variety of sources. Their main role in the system is to provide a shared vocabulary for describing things, but their formal definitions also permits some reasoning over concepts.³

Compendium

Compendium (Buckingham Shum et al., 2006) is a concept-mapping tool to support real-time, collaborative sense-making by linking argumentation and information; and by providing a live document of this process, it also acts as a collective group memory. It uses and extends concepts introduced in the Issue-Based Information System (IBIS) and Questions-Options-Criteria approaches for collective decision-making. *Compendium* provides its user with one or more information spaces (or *maps*) into which *nodes*, each representing an idea, argument, document or some other entity, can be introduced and linked to other, related nodes. A hypertext paradigm allows the same node to appear on

² This response organisation is based on that devised for London. See <http://www.leslp.gov.uk> and (LESLP, 2004).

³ See <http://www.e-response.org/ontology/> for examples of the ontologies used.

different maps, allows a map to appear as a node on another map, and allows external information (documents, images, videos) to be referenced and accessed in context. This produces not ‘flat’, independent visualisations, but rather a rich set of cross-referenced views of the nodes and their relationships, reflecting more faithfully the non-linear nature of complex decision-making discussions and constituting a more evocative document of the process. Compendium provides the principal user interface to our e-Response system; hence, system use is focused upon making sense of the emergency and developing the response, with the other tools invoked to help as necessary. To do this effectively, Compendium was modified to allow simple context-sensitive, menu-driven invocations of the other tools. In addition, to ‘situate’ the sense-making specific node types (representing incidents, medical centres, etc.) are provided, and a street map of the affected area of London is used as a background image to the main map; as nodes are added to the conceptual map they are superimposed where they ‘occur’ on this cartographic plane.

I-X

I-X (Tate, Dalton and Stader, 2002) is a generic system architecture (and accompanying tool-suite) whose function is to support processes that develop ‘products’; in the context of emergency response a product might be (the creation and enactment of) a response plan. An I-X system consists of some federation of communicating human and computer agents, united by the <I-N-C-A> framework: communications are in terms of *issues* (unresolved problems), *nodes* (activities), *constraints* (which describe the world state) and *annotations* (meta-information about other items). I-X provides an underlying architecture for the e-Response virtual organisation, with the JESCC as one agent in this organisation: an I-X wrapper around Compendium, allows it to communicate with the other I-X agents – namely the Fire Brigade CSC, the Police CCC-IR, the Central Ambulance Control and the JESCC Liaison team. This enables the sense-making space of Compendium to be used to develop plans, which can then be relayed to other agents for enactment; in turn, these agents can raise issues to the JESCC for deliberation (Tate *et al.*, 2006).

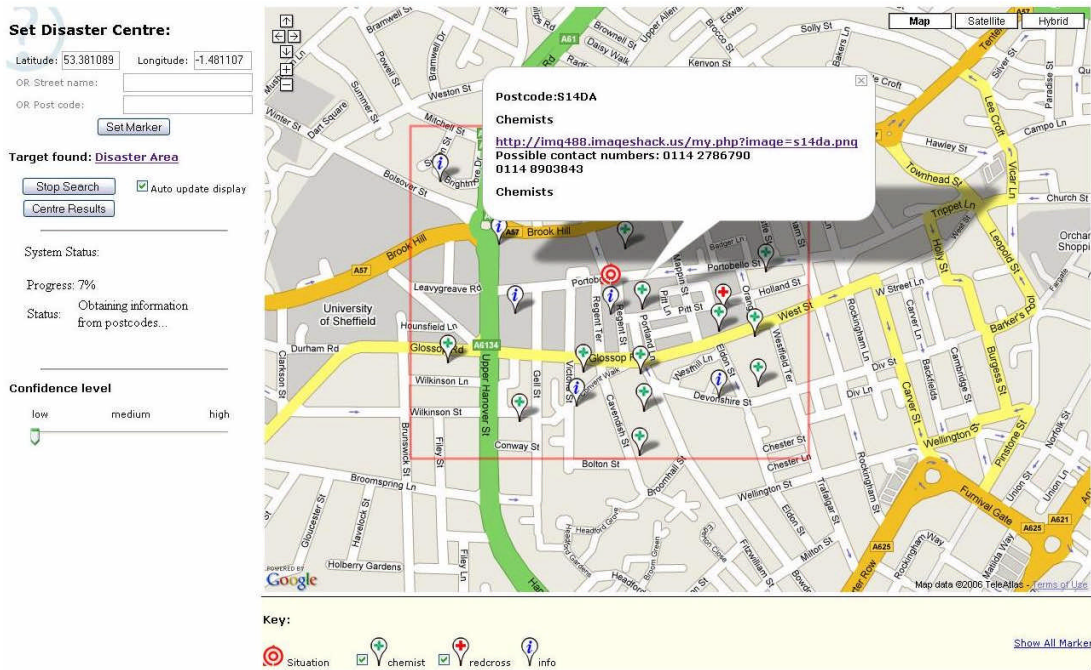


Figure 1. Armadillo e-Response knowledge-base interface

Armadillo

For any but the most immediate of emergency requirements, the task of creating and maintaining repositories of potentially useful resources does not scale. *Armadillo e-Response* (Chapman and Ciravegna, 2006) introduces a methodology to address this by constructing in real time a repository of local information. It comprises a number of service components that concurrently collate increasingly refined knowledge of resources in the vicinity of the emergency. Cross-referenced information about postcodes and geospatial coordinates triggers information retrieval using web search engines and web crawlers. The retrieved resources must be indexed; Armadillo aids this by applying a pre-trained classifier to the resources to classify them into categories of interest such as possible triage

locations, marshalling zones, etc. Next, natural language processing (NLP) techniques are used to extract contact details. Finally, the knowledge is stored in a triplestore. A map-based visual interface (Figure 1) allows exploration of the growing knowledge base.

CROSI Mapping System

One difficulty for any system of heterogeneous agents and knowledge – such as our e-Response system – is that of ensuring semantic interoperability, that is, of resolving references to the same entities expressed using different ontologies. The *CROSI Mapping System* (Kalfoglou and Hu, 2005) includes a number of linguistic, structural and semantic matching algorithms for concept mapping situated in a modular architecture. In the e-Response system, CROSI is used in an innovative manner to ‘translate’ terms so they are understood by the user and other tools. The nature of this task is such that there is always some degree of imprecision involved; yet, as will be seen, alongside a human operator able to filter suggestions contextually, CROSI is an effective approach to the mapping problem.

Photocopain

Photocopain (Tuffield *et al.*, 2006a) supports the annotation of photographs, and hence, the management of image repositories. It proposes annotations based on a combination of contextual information captured by a ‘semantic logger’ infrastructure and content-based information generated by image analysis techniques. Inspired by the use of mobile phone images during recent real emergencies, in the e-Response system a service has been added to Photocopain to allow members of the public to upload their photographs and the generated location and event annotations to a triplestore, where they represent a dynamic source of knowledge of the emergency.

AKTive Media

AKTive Media (Chakravarthy *et al.*, 2006) uses semantic web and NLP technologies to help users annotate content across a variety of media. It provides an interface that supports ontology-based and free text annotation of text and images, and can assist the user by suggesting terms to describe content by, for instance, processing candidate annotations proposed by Photocopain. For the e-Response system, the ontology-based annotation of potentially useful images is assumed to have been done prior to the emergency, and the annotations uploaded to a triplestore.

OntoCoPI

OntoCoPI (Alani *et al.*, 2003) identifies communities-of-practice amongst ontology instances that describe humans, documents, projects, etc., by analysing the semantics of the relationships that occur between instances, and by applying network analysis techniques to assess the connectivity and cumulative density of instances. The analysis is done at run time, and can be activated with queries for ‘seed’ instances and concepts to explore the different communities that exist around these. For the e-Response system, a knowledge base of civil engineers was generated from lists maintained by professional bodies and described according to appropriate ontologies, and then queried via an HTTP request, sent to the OntoCoPI service.

The Commitment Management System

The role of the *Commitment Management System* (CMS) in the e-Response system is to recommend ‘good’ allocations of available emergency service resources where those resources are over-constrained and the demand outweighs the number available. With commitments on resources represented as constraints, and each given a relative ‘importance’ utility value, CMS searches for partial solutions where only a subset of constraints is applied. The choice of which to apply depends on the utility metric used in conjunction with the utility values; and these also allow each potential solution to be given a ranking score based on the constraints it satisfies (Chalmers *et al.*, 2004). Communication with CMS is via a constraint interchange format for the semantic web called *CIF/SWRL* extended with an ontology for modelling soft constraint satisfaction problems (Preece *et al.*, 2006).

E-RESPONSE SYSTEM: DEMONSTRATION

The JESCC has been convened, and communications established and tested. The JESCC computer support officer starts Compendium and loads a street map of the affected area of London. Then the first message arrives....

10:00:32 Message from Police CCC-IR: “Major fire at Fenchurch St. Station. London Fire Brigade notified”

From this cue, the support officer begins to populate the conceptual map of the London emergency by introducing a new *incident* node, placing it over Fenchurch St. Station on the street map, and then converting the message into the text of an *idea* node (used generally to represent alerts and information). This node is tagged automatically with the name of the sending agent and the time it was received. The officer relates the message node to the incident node by dragging a link from the former to the latter. Prompted by the JESCC members around him, the officer adds and links a *question* node with the text “scale of incident?” – this represents an issue that the JESCC wants resolved if it is to assess fully the situation. Right-clicking on this question node gives a menu of options: the type of node means these include sending the node text as an I-X issue to the other I-X agents in the system – and in this case, the officer sends it to the *Police-CCC-IR* agent. Figure 2 shows the state of the Compendium interface after handling this first message. It can be seen that the support officer has begun to construct a conceptual map of the incident by adding and relating nodes overlaid on a cartographic map to give some indication where the information, activities and issues are ‘occurring’. He will continue to develop the map in this way for the duration of the emergency, adding and linking established facts, incorporating relevant information, recording decisions made in the JESCC and so on.

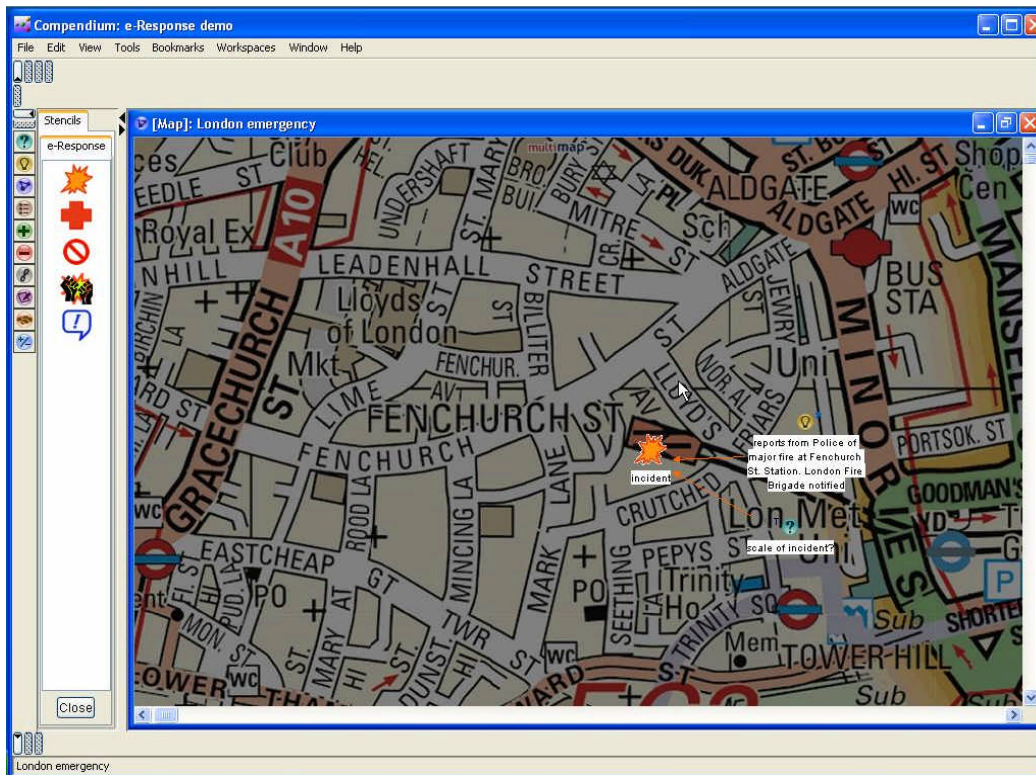


Figure 2. The visual state of the system after the 10:00:32 message has been handled.

10:02:42 Notification from Fire-Brigade-CSC of resource allocation to Fenchurch St. Station Fire

The next message to be received is generated by the Fire-Brigade-CSC I-X agent (Figure 3), and contains details of the allocation of resources to the station fire. This information is described using a structured XML template. The support officer loads the details into a new hyperlinked ‘response event’ map within the current working map; the *map* node representing this new map is linked to the corresponding incident node. The received information includes the (lat-long) location of the response, the quantity of resources (here each resource is assumed to be a fire engine) the event requires, and the estimated time that the response will last (Figure 4). The location information allows Armadillo to be invoked (again via a context-sensitive pop-up menu) to search for potentially useful resources in the vicinity.

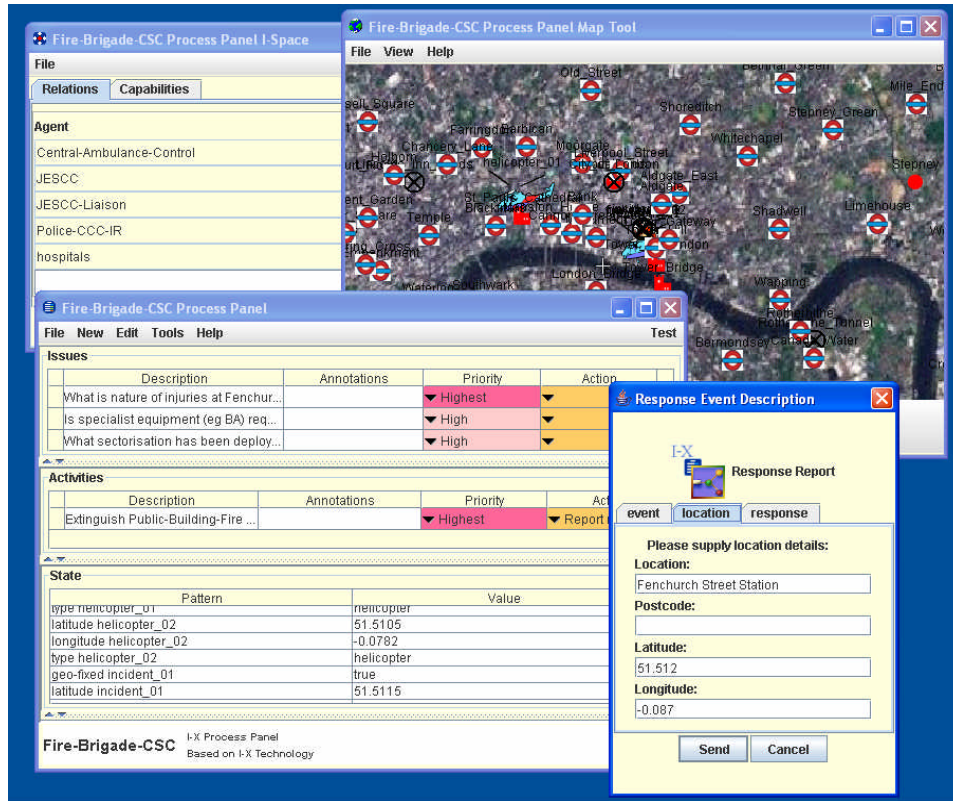


Figure 3. Using an I-X Panel to report the response to the Fenchurch St. Station fire to the JESCC. The main panel shows current activities and issues; other tools provide state visualisation and knowledge management support.

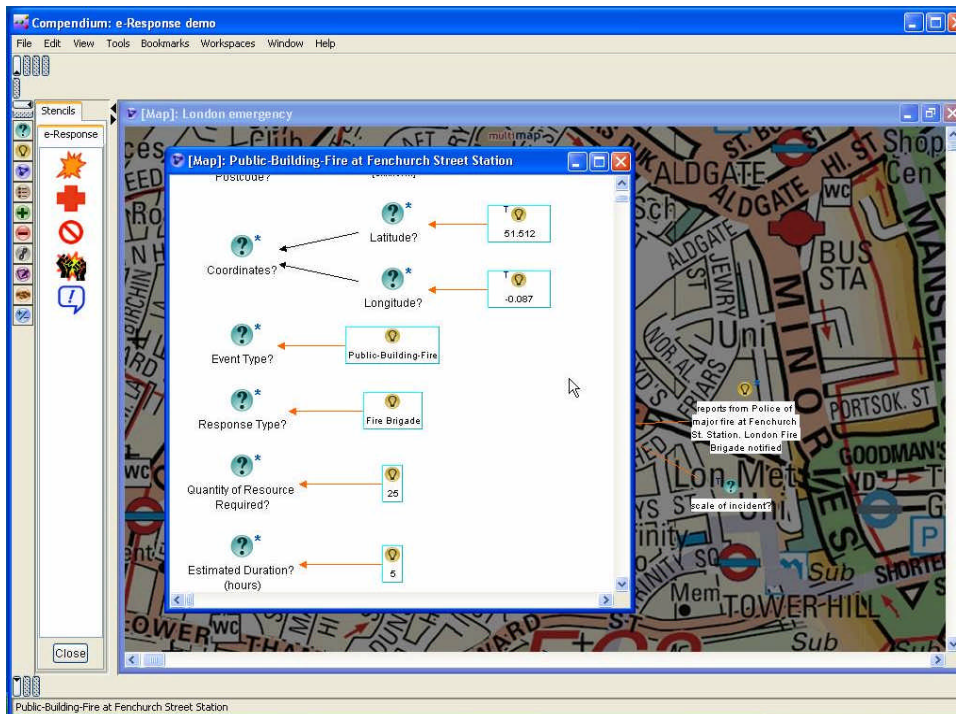


Figure 4. Some of the information in the 'resource allocation map' for the Fenchurch Street Station incident.

10:05:12 Message from JESCC-Liaison: “Message from fire brigade officer at Fenchurch St. Station: ACFO has deployed SSUs at the scene”

This message, relayed from a fire-fighter on the ground, is couched in acronyms and somewhat cryptic. To make sense of the message and handle it appropriately, the support officer must first understand it; he will attempt this using CROSI. He first adds the node to the map, and links it to the incident. Next he ‘transcludes’ the node: this operation creates a new map containing the node, which provides him with a working space. He introduces question nodes for each of the two terms he does not understand: “ACFO” and “SSU”. Now, right-clicking on the first of these, he invokes CROSI with a fire brigade ontology (since the message source suggests the term relates to the Fire Brigade). CROSI returns several options, some more plausible than others (Figure 5); the officer is able to filter the less appropriate ones and select that which makes most sense in context – “AssistantChiefFireOfficer”. He now invokes CROSI again, this time for “SSU”, and is told that this maps to, amongst others “ScientificSupportUnit”, which seems to make sense here. Hence, he has now ‘deciphered’ the message, and can see that it is notification of the deployment of a specialised unit, and that, beyond noting this, no further action is required of the JESCC.

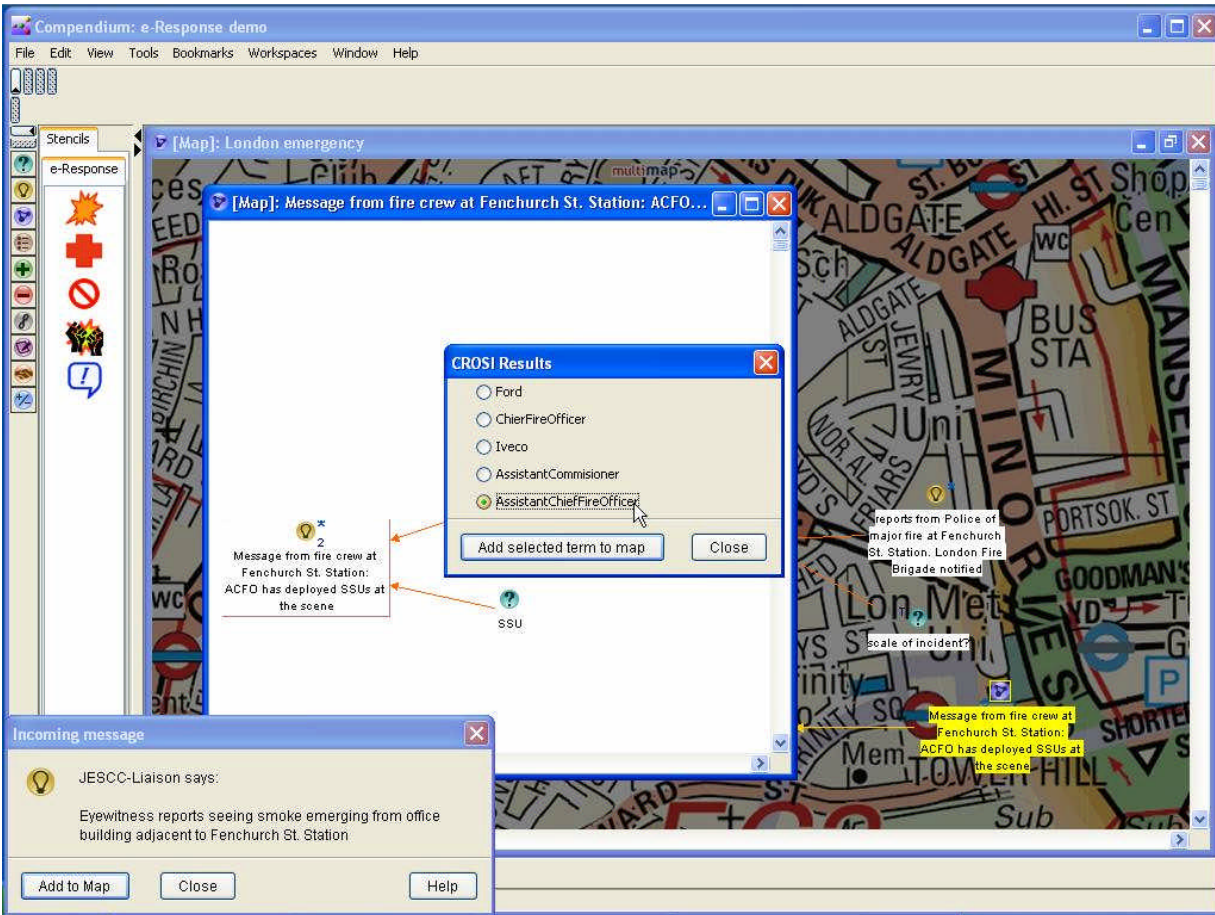


Figure 5. Decoding jargon using CROSI.

10:08:40 Message from Police-CCC-IR: “Request details of burns care units in Fenchurch St. Station vicinity”

Once the message has been added to the map as a question node, the support officer turns to his web browser to look for an answer from Armadillo, which now shows the closest known burns care facilities are to be found at St. Mary’s, Paddington, information which is added and relayed to the Police-CCC-IR.

10:09:43 Message from JESCC-Liaison: “BBC showing mobile phone images of London fires”

This message, originating from a member of the JESCC liaison team detailed to monitor media coverage, is the cue for the support officer to view, in a browser, the Photocopain image upload site. Among the images are some that

appear to be of the London emergency; accordingly he starts the Photocopain analysis. This suggests that the images show an incident at Lloyd's Building, about a quarter of a mile from the Fenchurch St. Station fire (Figure 6).

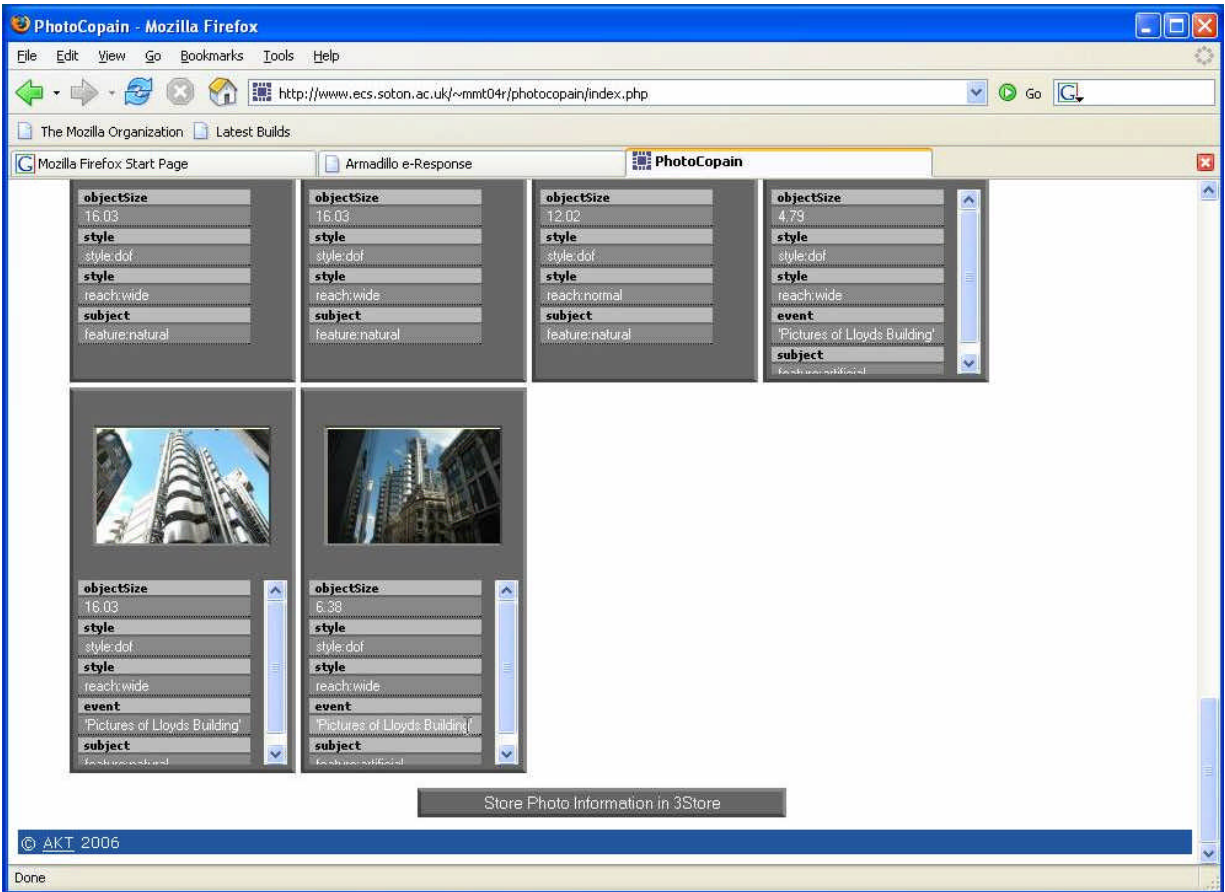


Figure 6. The results of Photocopain photograph analysis.

10:10:38 Notification from Fire-Brigade-CSC of resource allocation to Lloyd's Building Fire

10:15:04 Notification from Fire-Brigade-CSC of resource allocation to Lower Thames St. Warehouse Fire

These messages are notifications of further allocations of fire brigade resources, the first confirming the Photocopain analysis. As before, for each a new incident node is added to the map, and the details of the response imported.

10:18:55 Message from Fire-Brigade-CSC: "request details of upper-storey exit routes from Lloyd's Building"

With people apparently trapped, the Fire-Brigade-CSC asks for exit routes from the upper levels of the building. This prompts the support officer to query a triplestore for images related to the location, via a menu item available within the Lloyd's Building fire response map, since this contains lat-long coordinates. This generates and sends a request to the triplestore, which returns the URIs of a number of images. A *reference* node is automatically created for each URI, causing thumbnails of the images to be added (Figure 7). Among these are plans and photographs of Lloyd's Building; in addition to location meta-information, some have features annotated (using AKTive Media) with terms from a building ontology. Double-clicking on a node loads the annotated image into a browser where the annotations can be seen (Figure 8). The URLs of appropriate images are then sent to the Fire-Brigade-CSC.

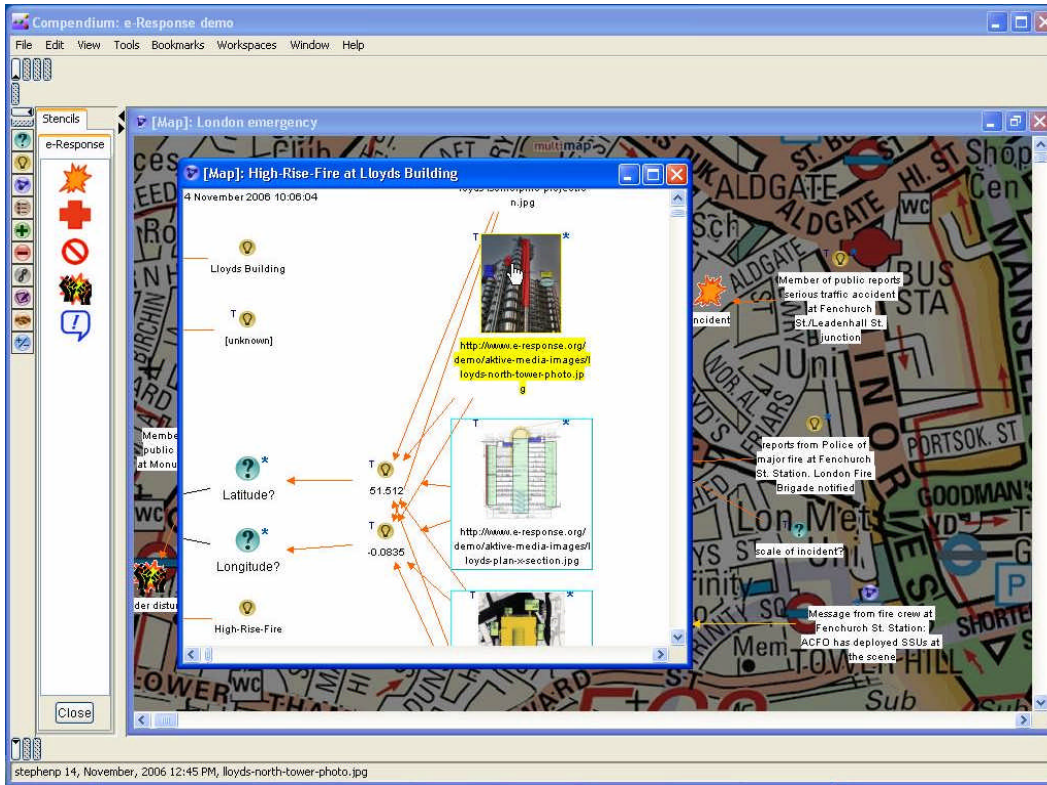


Figure 7. Thumbnail images of the Lloyd’s Building location retrieved from the 3Store.

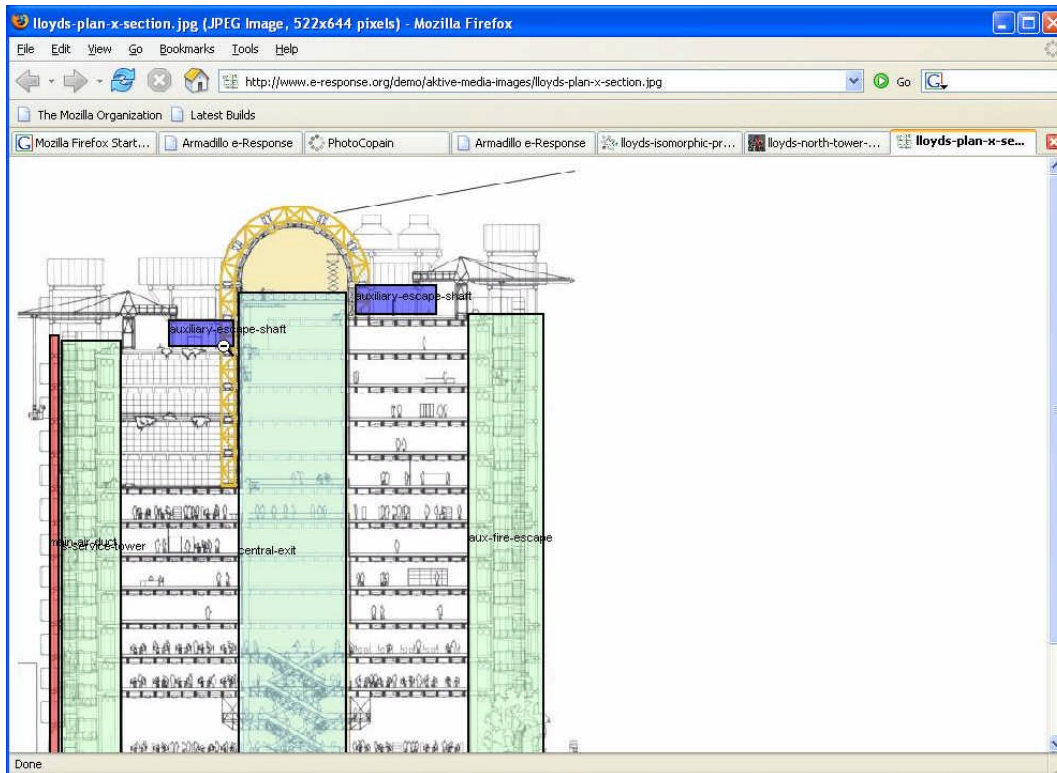


Figure 8. Lloyd’s Building image annotated using the AKTive Media tool.

10:24:03 Message from Fire-Brigade-CSC: “brigade engineers at Fenchurch St. Station report incidence of creeps in building fabric”

This is another cryptic message using jargon: again, CROSI will be used to try to make sense of it. As earlier, the message is added as a node, linked to the Fenchurch St. Station incident, and transcluded to a new working map. The support officer invokes CROSI for the term “creep” and with a building ontology, since, from the message context, he infers that the term is related to buildings and structures. Among the mappings suggested by CROSI is one concept, “ConcreteCrack”, that seems to fit the context, and so he adds this into the map. Since the idea of cracks in the Station building sounds serious, the JESCC decides that expert advice is required. A menu option for the new concept node allows a triplestore to be queried for fields of practice related to that concept. This returns a list of related fields. The officer now chooses a relevant-sounding entry from this list, “Concrete-Structures”, and adds it as a new node. With this new node he can query OntoCoPI for experts clustered around this seed field. OntoCoPI returns half-a-dozen experts, and further calls to a triplestore retrieve contact details, including home web-pages, for each. The web-page of one of these experts says he is based in London and, moreover, that he has experience in railway engineering. On the telephone, the expert says that, in his opinion, creeps in the Station fabric do not pose an immediate threat and can be ignored for the time being, advice which is accepted by the JESCC. A node representing this decision is added (Figure 9), and then relayed, as an I-X activity, to the Fire-Brigade-CSC.

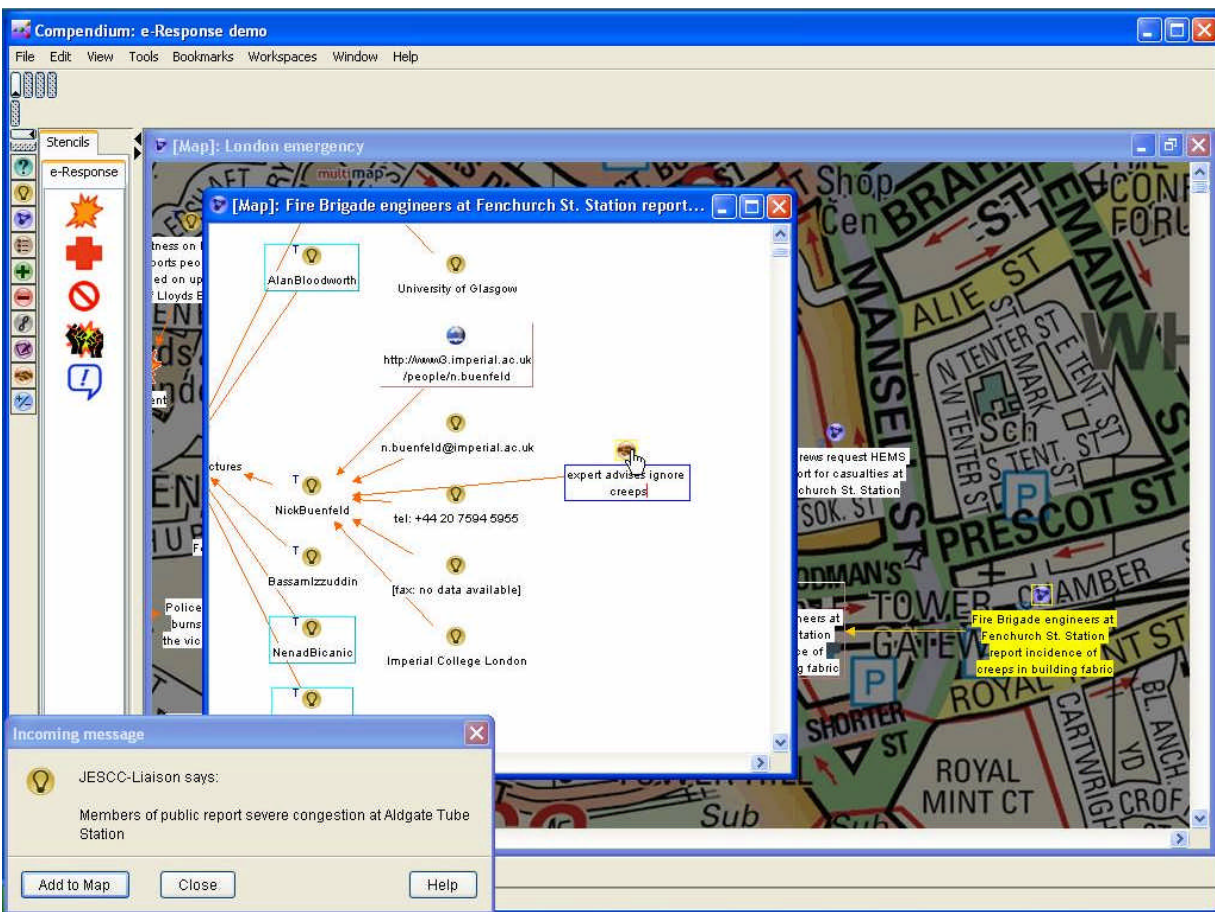


Figure 9. Using expert advice to resolve the matter of creeps.

10:33:21 Request from Fire-Brigade-CSC for resource allocation to Tower of London Fire

This message refers to another new incident, a fire at the Tower of London. Although structurally identical to previous notifications, there is one important difference: whereas the previous messages indicated that resources were being deployed at the incident in question, this message is a *request* for resources. In other words, the desired response to this new incident demands more than the number of unallocated fire engines currently available, and as a result the Fire-Brigade-CSC has requested tactical guidance from the JESCC. To aid them in this the members of the

JESCC request the use of CMS to suggest allocations. Accordingly the support officer, once he has recorded the new incident on the map, copies the existing allocations into a new resource management working space, and attaches relative priorities to each of the current incidents: the Fenchurch St. Station fire is accorded a ‘high’ priority (there have already been reports of casualties at what is usually a busy station); the Lloyd’s Building fire a ‘medium’ priority (no reports of casualties); the Lower Thames St. fire a ‘low’ priority (it is thought that few people will be in immediate danger there). The new incident at the Tower of London is given the ‘highest’ priority – the Tower is likely to be busy with tourists and staff, and, furthermore, is of symbolic importance for the city. CMS is invoked with information about the resource allocations/requests, durations and priorities for the incidents and the total number of available fire engines. The CMS response is to suggest that resources are diverted from the Lloyd’s Building and Lower Thames St. fires to the Tower of London. The JESCC members accept this advice, which the support officer records as a decision node (Figure 10) that he then relays as an I-X activity to the Fire-Brigade-CSC.

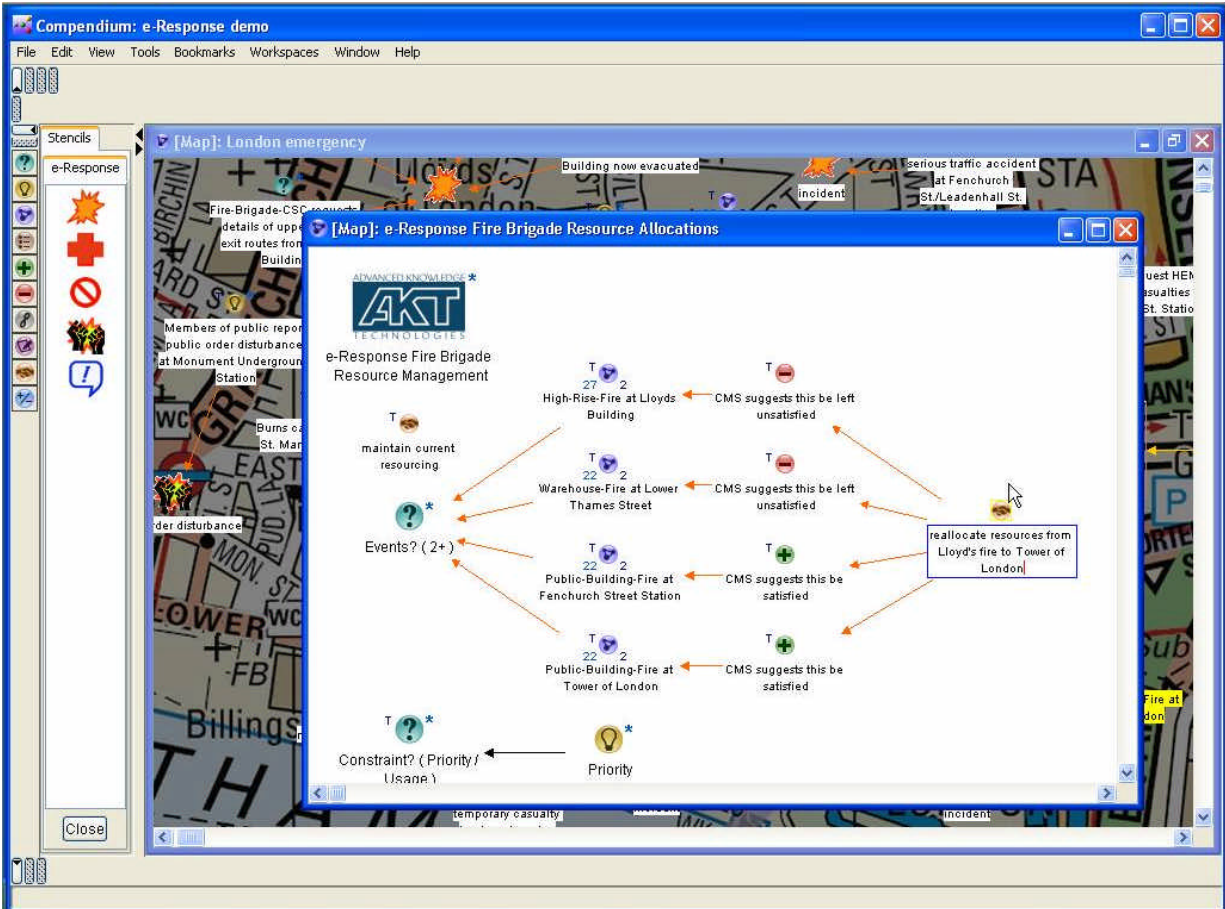


Figure 10. Accepting the CMS resource reallocation.

SUMMARY

The above demonstration illustrates how a loosely configured system based on underlying AI and semantic web techniques, and underpinned by ontologies and triplestore knowledge bases, can support emergency response:

- By providing an ‘intelligence space’, encompassing not only sense-making, but also information-mapping, collective memory capture, and decision-making support facilities for the group (Compendium);
- By situating the sense- and decision-making in a framework that encourages a structured and principled approach to shared activity across a virtual organisation of agents (I-X);
- By providing access to a concept mapping, disambiguation and ‘deciphering’ service (CROSI);
- By acting as a conduit for dynamic information gathering, and the visualisation of the results (Armadillo);

- By tapping into alternative communication channels, and analysing their content (Photocopain);
- By enabling the semantic indexing and retrieval of graphical information (AKTive Media);
- By helping to identify and access expertise within (and beyond) the response organisation (OntoCoPI);
- By applying computation techniques to resource management problems (CMS).

CONCLUSIONS

This paper discusses a system that is composed of and constructed using a number of different AI/semantic web technologies for the purpose of supporting the high-level tactical command of the response to a large-scale emergency. More specifically, the system is structured around a sense-making tool representing one agent within a wider collaborative enactment framework, based on the definition of the form, roles and responsibilities of an existing major emergency response organisation, LESLP. While no evaluation of the system has been attempted and, in many respects, the semantic web is still in its infancy, we hope that this system and its description present a compelling case for the role that technologies of the sort discussed here could come to play in assisting in dynamic situations where access to the 'right' information is crucial.

ACKNOWLEDGMENTS

This work is supported by the Advanced Knowledge Technologies IRC project, sponsored by the UK EPSRC under grant number GR/N15764/01. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing official policies or endorsements, either expressed or implied, of EPSRC or any other member of the AKT IRC. The maps shown in figures 2, 4, 5, 7, 9 and 10 are © Collins Bartholomew Ltd; the authors gratefully acknowledge the permission of the copyright holder to reproduce these maps in this paper.

REFERENCES

1. Alani, H., Dasmahapatra, S., O'Hara, K. and Shadbolt, N. (2003) Ontocopi - using ontology-based network analysis to identify communities of practice. *IEEE Intelligent Systems*, 18(2), 18-25.
2. Berners-Lee, T., Hendler, J. and Lassila, O. (2001) The semantic web, *Scientific American*, May 2001, 35-43.
3. Buckingham Shum, S., Selvin, A., Sierhuis, M., Conklin, J., Haley, C. and Nuseibeh, B. (2006) Hypermedia Support for Argumentation-Based Rationale: 15 Years on from gIBIS and QOC. In: *Rationale Management in Software Engineering*, A.H. Dutoit, et al. (eds), Springer-Verlag: Berlin.
4. Chakravarthy, A., Ciravegna, F. and Lanfranchi, V. (2006) Cross-Media document annotation and enrichment, In, *Proc. 1st Semantic Web Authoring and Annotation Workshop (SAAW2006)*, Athens, GA., Nov 2006.
5. Chalmers, S., Preece, A., Norman, T. and Gray, P. (2004) Commitment Management Through Constraint Reification, In, *Proc. 3rd Int. Joint Conf. on Autonomous Agents and Multi-Agent Systems (AAMAS 2004)*, pp. 430-437, IEEE Press, New York, USA.
6. Chapman, S. and Ciravegna, F. (2006) Focused Data Mining for Decision Support in Emergency Response Scenarios. In, *Proc. Int. Semantic Web Conference 2006 Workshop on Web Content Mining with Human Language Technologies*, Athens, GA., USA, November 2006.
7. Harris, S. (2005) SPARQL Query Processing with Conventional Relational Database Systems, WISE Workshops 2005, pp. 235-244.
8. HMFSI (2002) Fire Service Manual, Volume 2 Fire Service Operations, Incident Command, HM Fire Services Inspectorate Publications Section, London: The Stationary Office. Crown Copyright.
9. Kalfoglou, Y. and Hu, B. (2005) CMS: CROSI Mapping System – Results of the 2005 Ontology Alignment Contest}, In, *Proc. K-Cap'05 Integrating Ontologies Workshop*, Banff, Canada, Oct 2005, pp. 77-84.
10. LESLP (2004). *Major Incident Procedure Manual*. 6th Ed, July 2004. Available at: <http://www.leslp.gov.uk>
11. Preece A., Chalmers, S., McKenzie, C., Pan, J. and Gray, P. (2006) Handling Soft Constraints in the Semantic Web Architecture. In, *Proc. WWW 2006 Workshop on Reasoning on the Web (Row 2006)*, Edinburgh, UK.
12. Tate, A., Buckingham Shum, S.J., Dalton, J., Mancini, C. and Selvin, A.M. (2006) Co-OPR: Design and Evaluation of Collaborative Sensemaking and Planning Tools for Personnel Recovery, Open University Knowledge Media Institute, Technical Report KMI-06-07, March 2006.

13. Tate, A., Dalton, J. and Stader, J. (2002) I-P² – Intelligent Process Panels to Support Coalition Operations. In *Proc. 2nd Int. Conf. on Knowledge Systems for Coalition Operations (KSCO-2002)*. Toulouse, France, 2002.
14. Tuffield, M.M., Harris, S., Dupplaw, D.P., Chakravarthy, A., Brewster, C., Gibbins, N., O'Hara, K., Ciravegna, F., Sleeman, D., Shadbolt, N.R. and Wilks Y. (2006a) Image Annotation with Photocopain, In, *Proc. Semantic Web Annotation of Multimedia (SWAMM-06) Workshop at the World Wide Web Conference 06*, May 2006.