

Using Agents in Virtual Environments to Assist Controllers to Manage Multiple Assets

BEER, Martin <<http://orcid.org/0000-0001-5368-6550>>, ALBOUL, Lyuba <<http://orcid.org/0000-0001-9605-7228>>, NORLING, Emma and WALLIS, Peter

Available from Sheffield Hallam University Research Archive (SHURA) at:

<http://shura.shu.ac.uk/5661/>

This document is the author deposited version. You are advised to consult the publisher's version if you wish to cite from it.

Published version

BEER, Martin, ALBOUL, Lyuba, NORLING, Emma and WALLIS, Peter (2012). Using Agents in Virtual Environments to Assist Controllers to Manage Multiple Assets. In: Workshop on Cognitive Agents for Virtual Environments, AAMAS 2012, Valencia, Spain, 4th June 2012.

Copyright and re-use policy

See <http://shura.shu.ac.uk/information.html>

Using Agents in Virtual Environments to Assist Controllers to Manage Multiple Assets

Martin D. Beer¹, Lyuba Alboul², Emma Norling³, and Peter Wallis⁴

¹ Communication & Computing Research Centre
Faculty of Arts, Computing, Engineering & Sciences
Sheffield Hallam University
Sheffield, United Kingdom
m.beer@shu.ac.uk

² Centre for Automation & Robotics Research
Faculty of Arts, Computing, Engineering & Sciences
Sheffield Hallam University
Sheffield, United Kingdom
LAlboul@shu.ac.uk

³ Centre for Policy Modelling,
Manchester Metropolitan University
Manchester, United Kingdom
norling@acm.org

⁴ The NLP Group, Dept of Computer Science, University of Sheffield
Sheffield, United Kingdom
p.wallis@dcs.shef.ac.uk

Abstract. Search and rescue operations often require complex coordination of a range of resources, including human and robotic resources. This paper discusses a proposed new framework that allows agent technology to be used in conjunction with a virtual environment to provide a human controller with an effective visualisation of the distribution of a collection of autonomous objects, in our case, Unmanned Aerial Vehicles (UAVs) so that they can be managed in a way that allows them to successfully complete the task in the minimum possible time. It is our contention that to do this effectively there needs to be two-way initiation of verbal conversations, but that it is not necessary for the system to completely understand the conversations required. An example scenario is presented that illustrates how such a system would be used in practice, illustrating how a single human can verbally communicate with a swarm of semi-autonomous actors verbally and envisage their activities in a swarm based on the visual cues provided within the virtual environment. An agent-based solution is proposed that meets the requirements and provides an effective command station that can manage a search using a collection of UAVs effectively.

1 Introduction

Agent technology has been used extensively with virtual environments for a range of educational [1], gaming [2] and training applications [3]. This paper considers another situation where virtual environments could provide a vital link with between humans

and real world activities. One area of considerable interest is the use of multiple robots or other autonomous agents to perform some large scale cooperative task, such as search and rescue [4]. Robots are considered to be a valuable asset in search and rescue activities as they can be sent into areas which have not been made safe for human rescuers. Initially single robots were used with a human controller. While it is relatively straightforward for a single human to remotely control a single robot, teams of robots can be used more complex and efficient searches. It has still proved necessary for individual human operators to provide both the intelligence and coordination [5].

This paper proposes a framework that provides a single operator with the means to control large teams of autonomous agents in complex operations via a virtual environment. The human operator is immersed in a virtual environment where the (semi) autonomous physical agents are represented by virtual agents. The operator coordinates the behaviour of physical agents by interacting with their counterparts in the virtual representation of the physical world. It is an open question as to where the intelligence lies but typically the physical agents will have limited (if any) cognitive abilities, that simply enable the physical agent to operate autonomously for short periods of time. The on-board processing limitations of these platforms precludes much in the way of higher-level reasoning; instead, this would be performed by the physical agent's virtual counterpart, who would receive sensory data from the physical agent and send instructions back to it. In addition, these virtual agents are also "embodied conversational agents," allowing the human operator to coordinate the agents through a natural language interface. The use of a spoken language interface in this scenario has two advantages. The first that language is our best example of a mixed initiative interaction in which the human can initiate an interaction by issuing a command or requesting information, but the agent can also provide information in a timely manner without being asked. The second advantage is that natural language allows us humans at least to negotiate new information.

The scenario we use to illustrate the problems involved is of a search and rescue operation involving unmanned aerial vehicles (UAVs). A single human controlling (say) twenty surveillance UAVs introduces a range of problems but, for such mixed teams to work at all, a certain level of autonomy for the UAVs is required. The assumption is that each has sufficient intelligence that once it has received specific instructions, it is able to maintain itself on station and to perform its allotted task, which may include relaying messages to peers over the horizon from the base station, and so out of direct contact, acting as a mobile telephone relay station and collecting and relaying sensor data. Sensors may be a standard set that are common to all individuals, or they may be configured specifically for each operation. For the purposes of this analysis it does not matter. It also requires a means of communication – both machine with machine and machine with human. We therefore provide a virtual world for the human operator in which real UAVs are represented by embodied conversational agents (ECA) that can be seen performing their duties, and that can be conversed with in plain English. These agents have limited cognitive abilities and these limitations are, like the plans and goals of each agent, something that the agent can talk about. A spoken language interface is not only intuitive and flexible, being able to "call out" allows agents to initiate a conversation in a graded manner that is difficult with simple alarms. The challenge

however is to provide an agent that can hold even simple conversations. From the very early days of AI research it has been possible to hold a conversation with a machine in a limited domain [6, 7]; the problem is that we humans are not very good at sticking to the topic. This issue has been considered [8, 9] together with ways to bring the user back on topic without them noticing.

The approach taken in this paper is to develop a typical scenario based around real situations that rescue services experience, and to construct a suitable system to meet that need. The scenario chosen is close at hand, and therefore rescue personnel with real experience of similar situations are available to advise us. This approach allows the system to be validated with these real users at an early stage, and they can be updated and comment on each development cycle. This leads to a more usable and functional system as all the components are validated at each stage. This is particularly important with the development of the “virtual world” as this needs to be functional enough to represent the actual situation on the ground, but abstract enough to show the essential information without distracting the controller with unnecessary detail.

2 The Scenario

The scenario is based on a search and rescue mission in the Northern Peak District in the United Kingdom, which despite the relative smallness of the area covered, involves most, if not all, of the activities that such operations typically require. This is because of the poor communications across the area, its relative popularity particularly with inexperienced and ill-equipped visitors who have easy access from nearby large cities, and the rapid changes in weather, particularly in winter, when time is often of the essence in successfully evacuating casualties. BBC reports [10–14] illustrate the wide range of incidents that the rescue services have to deal with within this small area.

2.1 A Typical Incident

The scenario starts with a set of UAVs each carrying a mobile phone repeater station “brick” as used to fix dead-spots in the mobile telephone network, for example within steep valleys (locally called cloughs). Each UAV flies at a fixed height and, within limits, can re-position where it “loiters” based on the amount of communication traffic it is relaying and how close it is to other UAVs in the team. The point however is the model for human-machine interaction based on the virtual world and conversational representations of the robots.

Sergeant Jones is three hours into a shift providing radio coverage for a search and rescue mission over the Peak District. A Duke of Edinburgh Award Group set off from Edale to follow the southern part of the Pennine Way across Kinder Scout. They have failed to report in as expected, the weather is closing in and a search and rescue operation has been launched. They are not responding to mobile phone messages, but communication in the area is notoriously difficult.

The Kinder plateau is a large upland area with few markers, and it is easy to become disorientated in poor weather. The fear is that the group may have become lost, possibly separated, and that as the weather deteriorates towards nightfall they will become

increasingly at risk both from the terrain, which includes peat bogs and deep ravines, and exposure. A ground search has commenced involving Mountain Rescue, the Fire and Rescue Service, National Trust and National Park Rangers. Figure 1 shows the area of the National Park. The area of the scenario consists of northern upland fells where there are few roads, and plenty of opportunities for hiking and other outdoor pursuits.



Fig. 1. A Map of the Peak District National Park. The Scenario considers the northern fells north of Edale where roads and habitations are very sparse.

Because of the diversity of groups involved, the primary means of communication is mobile telephones. In view of the poor communications coverage, Sergeant Jones' team has been asked initially to back up and "fill out" the existing mobile telephone infrastructure's coverage. Jones is the controller for twenty semi-autonomous UAVs that can provide temporary mobile telephone base stations. She is standing in a virtual environment provided by a data cave with images of The Peak topology displayed on all

four walls, as viewed from the centre of the search area at a height of 2km as shown in Figure 2. The system can modify the lighting to represent day and night, and can provide various shadings to assist Jones in both visualising conditions on the ground and have a clear overview of the topology over which she is working. She can also superimpose simulated weather patterns generated from Meteorological Office data to assist her in her task. She can move to a different location by sitting in a virtual electric wheelchair and motoring in any direction, but usually she stands or walks about. In either case, her effective position as represented in the cave moves to follow her. Looking west she can see, super imposed on a recent photograph of the weather in that direction, the flight paths of airliners heading to and from Manchester Airport (the flight paths cross the search area), and north of that she can see a set of icons representing the electrical storm coming in from the west. Also projected on her view are twenty coloured deltas, each representing the position and travel of one of her UAVs.



Fig. 2. The Control Station

3 Meeting the Requirement

She has just made a cup of tea when a voice off to her right calls out

“Sir”

She turns to it and addresses the blue delta

“Yes Blue?”

“Sir, I am low on fuel and request permission to return to base.”

“Granted”

The blue delta turns west and heads off to its temporary landing field for a service and to refuel. Jones knows that it will return in about 90 minutes and need redeploying. Ten minutes later she notices one of the deltas is heading north. She points at it and says

“You. What are you doing?”

A voice from that direction says

“Sir, there has been no radio activity in my area and I am following a lead.”

Jones asks to listen in on the call and hears a conversation about a party. That UAV is, it seems, following a car travelling westwards towards the Snake Pass (the A57) and away from the search area. She looks about and sees a fairly vacant area just west of the Howden Reservoir. She calls up the map reference for this area and directs the stray UAV saying:

“Head to 53 degrees 26 north by 01 degrees 42 west and look there.”

The UAV’s agent confirms by repeating the coordinates it thinks it heard:

“53 26 north by 10 42 west.”

which would take it a long way west and far away from the search area. Jones corrects it by saying:

“No no, zero one forty two west.”

The agent confirms that it now has the correct coordinates:

“Oh, 53 26 north by 01 42 west. Okay.”

and sends them to the UAV which takes up position as requested.

A bit later a new UAV appears from the west. One of the current UAV team – White – interrogates it (it is still outside of direct communication with the base station) and (White’s virtual agent) calls out

“Sir.”

Sergeant Jones says

“Yes White.”

“A Shadow AAI RQ-7 has come to join the search. Can you see it?”

Jones looks in the appropriate direction in the virtual world and sees a new delta that has just appeared. She interrogates it with:

“Yes, thank you” says Jones. She goes on: *“RQ-7, what are your capabilities?”*

The virtual agent for the new UAV pops up a photo of its physical agent with some detailed text and introduces itself by saying:

“I am a Shadow RQ-7 with a gimbal-mounted EO/IR camera, standard communications monitoring a G23 PA system and about 6 hours of fuel.”

Jones does not know the capabilities of the PA system so asks the new UAV’s agent to explain:

“How do you use the PA system?”

and is given this answer:

“I can fly at 200 feet and play a pre-recorded message. If the message is less than ...”

This gives her enough information for now and she closes the conversation with:

“Okay RQ-7”

Some time later Sergeant Jones decides to use the PA system to try to send a message to the lost hikers and gives the following instruction:

“go to 53 19 north, 10 42 west and then use your PA on a run to 53 23 north, 10 42 west with the following message: ...”

An hour before sunset a mobile phone message is received from the group, who are completely lost. The initial message is only relayed by UAV White, so there is insufficient information to pinpoint their location. Jones instructs the nearest UAVs (Red and Indigo) to form a holding pattern around White and to listen for further messages. She then calls the number from which the original message was sent, and the message is relayed by all the UAVs. When the call is answered, the location of the caller is pinpointed, and Jones hears that the group is together and safe. Jones relays this information to the rescue teams and a search group with all-terrain vehicles is redirected to locate and evacuate them. As soon as it is clear that the entire group have been located, the other search teams are instructed to stand down and return to base. The UAVs are however still required to provide communication and track the search teams until they are all safely off the moor. As the weather closes in the live images are updated and visibility is effectively zero. Met Office data is however imposed on Jones' view allowing her to see the extent of the bad weather. One rescue team will be caught before they can reach safety and a UAV is tasked to stay with them. As the storm approaches, its virtual agent calls out

"Sir."

Jones replies

"Yes Red"

The virtual agent requests, based on the sensor information communicated from the UAV

"The weather is getting rough. Requesting permission to return to base."

Jones replies that she wishes the UAV to stay on station despite the weather

"No Red, stay with that signal."

A little while later UAV Red requests to return to base and is again denied. A while later Red reports that it will need to land in the next 15 minutes and again it is told to say where it is. Finally the fuel runs out and it requests coordinates for a crash landing. Jones instructs it to go down in Ladybower Reservoir. All the people are out and safe two hours after sundown with the loss of one UAV that is later recovered and refurbished ready for the next incident.

While the use of UAVs as described is a future vision, the situation in this scenario is typical of a wide range of deployments that the emergency services have to undertake in often hostile conditions. In this case, we are looking at comparatively large tracts of open country, rather than the more normal firefighting situation of searching in the confined spaces of burning buildings as was the case in GUARDIANS [15]. This in itself provides a different set of challenges. UAVs are seen as a useful tool both to search larger areas very quickly and to provide temporary communication when necessary in what are inevitably poorly covered areas. They are becoming increasingly autonomous, and no longer need individual 'pilots' to control them, but pilots previously were able to communicate naturally with an overall controller, providing an important coordination element. Our hypothesis is that humans are better able to coordinate complex sets of operations involving multiple participants by giving instructions and receiving feedback by voice, rather than text- or image-based interfaces that are more commonly used

when coordinating with robots. This scenario describes a mechanism that seamlessly integrates machine decision making and human interaction. Our research aim is therefore to explore its potential and identify its limitations. Within this project, we are developing the virtual world interface and a suitable simulation mechanism for development and evaluation. We see no barriers to success regarding a base-line implementation; the question is primarily how far can we push it?

3.1 Approach to Language

The approach we are taking is to minimise the use of full understanding, and instead pay attention to the social roles and cues, to politeness and capturing the details of how people manage social relations with language [16]. Following Tomasello [17], we view language in use as intentional and cooperative; as long as the system can maintain its status as a social actor, the user will assume a cooperative intent. As an example consider how the dialogue in the scenario is not as smart as it may look. Consider what might happen if a person unfamiliar with the language capabilities of the virtual agent tried to interrogate the RQ-7 with questions such as

“Do you have a camera?”

to which the formal answer might be

“Yes,”

but the response can be, quite appropriately,

“I am a Shadow RQ-7 with a gimbal-mounted EO/IR camera, standard communications monitoring a G23 PA system and about 6 hours of fuel”.

With this second response, the user assuming a charitable intent is unlikely to realise that the system cannot answer the question, but rather assume that the system is taking the opportunity to help out with other information that might be relevant. By taking a situated approach to the dialogue, the point is to discuss the behaviours available to the system rather than its equipment. Rather than question answering, our vision for the system looks more like browsing.

The scenario makes extensive use of the autonomous control capabilities of modern UAVs. This allows them to operate within the operational area without continuous human intervention. They are able to fly to a specified location and to remain there by flying a pre-programmed circuit without further interaction. Not only can they act singly, but when a number are deployed they will establish a search pattern that efficiently covers the area in question, and if communication to the base station is impeded – for example by uneven terrain, or inclement weather patterns – they will arrange themselves in such a way that some of the group act as relays to the others so that communication is always maintained as shown in Figure 3.

The key objective therefore is to provide an environment for a single controller to maintain effective control of all aspects of the search so that the UAVs effectively become part of the team, and work in coordination with the other search parties, rather than as an independent resource. The controller is in verbal contact with the Emergency and Search services on the ground, so it is only natural that verbal communication should be used to manage the UAVs as well.

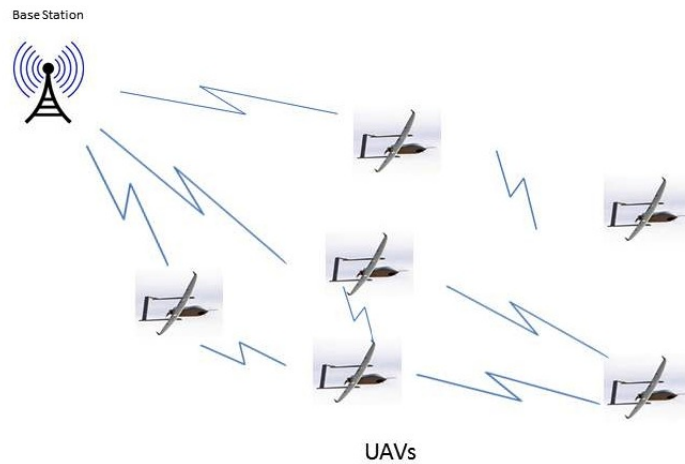


Fig. 3. The Communication Pattern of UAVs with the Ground Station

The UAVs are assumed to be standard production models, communicating with the system through special interface agents. The virtual counterpart of each UAV performs higher-level reasoning on its behalf, thus avoiding the need to modify the UAVs. The overall system architecture is shown in Figure 4.

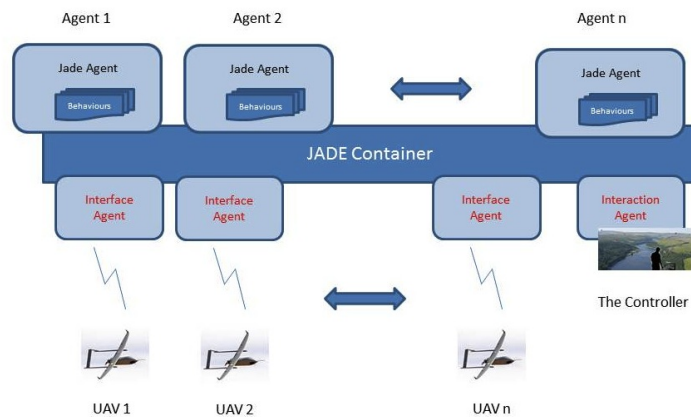


Fig. 4. The Agent Architecture to be used for Controlling a Group of Cooperating UAVs

From a language perspective, perhaps the first challenge is to decide on the appropriate level of representation to store explanations. At one extreme the person developing plans for UAV behaviour would provide an explanation of the behaviour in

text (say) and the text itself would be stored, and somehow stitched together when the agent is asked to explain its behaviour. At the other extreme, the provided explanations would be mapped into some formal representation of the text meaning and, when an explanation is required, the dialogue manager would use a text generation system would produce an appropriately tailored utterance. The second challenge is to have the dialogue manager choose what to say and when to say it in an appropriate manner – for example, which of the two responses to the “Do you have a camera?” question given previously would be appropriate?

3.2 The Agent Perspective

A layered approach has been adopted, as shown in Figure 5. The first layer provides the interface with the UAVs, and the last layer manages the virtual world and voice interfaces. It is the middle layer that is most interesting, as it provides the management and control at the core of the system. We propose to use BDI agents at this level, and an approach based on applied cognitive task analysis (ACTA) [18] to furnish these agents with domain-specific knowledge from experts [19].

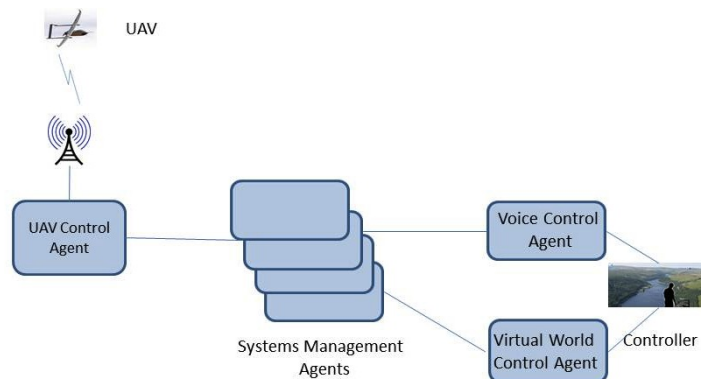


Fig. 5. Layering of Agents

A BDI architecture has been used on numerous occasions to drive simulated agents (both human and other) to drive robots and to drive dialogue. The GUARDIANS project has integrated a BDI language (JADE) with robot middleware (Player) and a simulation environment for search and rescue in a fire environment using the architecture shown in Figure 6[5].

3.3 The Use of BDI Architecture

The Belief, Desire and Intention (BDI) agent architecture was first developed in the 1980's at SRI International and implemented as PRS. [20] BDI is a software model de-

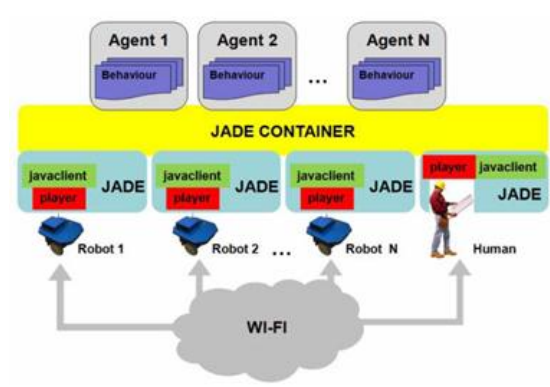


Fig. 6. The GUARDIANS Agent Architecture

veloped for programming intelligent agents. Superficially characterised by the implementation of an agent’s beliefs, desires and intentions, it actually uses these concepts to solve a particular problem in agent programming. In essence, it provides a mechanism for separating the activity of selecting a plan (from a plan library or an external planner application) from the execution of currently active plans. Consequently, BDI agents are able to balance the time spent on deliberating about plans (choosing what to do) and executing those plans (doing it). A third activity, creating the plans in the first place (planning), is not within the scope of the model, and is left to the system designer and programmer.

The idea is not to do planning, but to use plans to achieve goals. By separating plans from goals, the system can automatically re-plan on plan failure. Having recognised the problem it is easy enough for a programmer to write a rule that stops a UAV running out of fuel; the problem is there are so many commonsense rules and they are hard to notice until something goes wrong. Separating goals from plans goes some way toward alleviating the need to explicitly identify every rule. Here, the more important feature of a BDI approach is the way reasoning with BDI plans looks very much like the way we humans think other people think [21]. Based on our everyday “folk psychology,” a BDI approach provides a clear path to autonomous systems that can explain their behaviour [22].

The basis in folk psychology has another advantage: it can be used to facilitate knowledge acquisition [23, Ch. 6]. Following this approach, we hypothesise that we can develop cognitive agents to populate the virtual environment, using knowledge gathered from subject matter experts, with goals, plans and actions that are easily explicable to a human controller. Each of these cognitive agents will be linked to a single physical agent, performing the high-level reasoning and relaying the low-level commands to its physical counterpart.

First however, in order for a machine to perform any kind of interesting autonomous action, it must have sensor input. The classic is GPS positioning but this does not leave much room for autonomous decision making. Internal sensing (fuel/battery, vibration,

system functionality) are all sensor inputs that help, but some means of sensing the environment opens up the possibilities and make the nature of autonomous action more appealing and far more challenging. These sensor readings can be interpreted locally on board the UAV to provide local feedback control, or the information can be transferred back to the virtual agent for further processing. The question of *which* data should be transferred to the virtual agent is an open one: this agent must have sufficient data reason accurately about its physical counterpart, but will not necessarily require all the raw sensor data. There is also the possibility that the virtual agent may have incomplete knowledge about its physical counterpart, and it must be able to reason appropriately in this situation⁵.

3.4 The Virtual Environment

Development of the virtual environment starts with the construction of an initial 3D model of the scenario environment by using digital terrain models photographic images, bitmaps etc. A challenge will be to impose (integrate) the data obtained from the sensors (cameras, GPS) into this virtual world in real-time. At SHU some results in this direction have been achieved in the remits of the View-Finder project. The data from LRF (laser range finder) and camera have been fused to obtain a 3D photo-realistic representation of the environment, which then was 'inserted' in a 2D map obtained by a SLAM algorithm [24].

4 Further Work

We propose to take the knowledge acquisition process one step further by tagging plans for explanation. Behaviours can then be explained in terms of the plans driving those behaviours, and the conditions that caused those particular plans to be selected. The naturally hierarchical nature of BDI plans will facilitate the varying levels of detail that can be presented to the controller. It is the development of a suitable tagging mechanism for plans that will enable the generation of natural dialogue that will be the focus of this area of research.

The dialogue manager is an often-ignored part of a conversational system. For instance a classic approach to a natural language interface to a relational database is to treat the SQL query that would answer the query as the meaning of the text. The aim then is to translate the user input into SQL, perform the query, and then use text generation techniques to present the result. This is however not what people do. When we look at what people actually do, we find they put considerable effort into being polite [25], managing social relations [26] and, going out of their way to be helpful [17]. If the SQL query returns no matching records, it is not enough to say "no records match your request." People expect more. What do they expect, and how do our conversational partners go about providing it?

⁵ It should be noted though that a remote human controller of a single UAV may also find him-/her-self in this situation, and as such, our proposed knowledge acquisition approach should handle these cases in as much as they can be handled by a human controller.

There has been a growing interest in using partially observable Markov decision processes (POMDPs) [27] for spoken dialogue systems. The motivation is that Markov decision processes can model what the human will say next (based on some hidden brain state) and hence provide what the system should say next. Young at Cambridge has championed this approach in the UK and Lemon has been doing interesting work at Heriot-Watt [28]. Such machine learning techniques however are data hungry and simplifying assumptions need to be made. For instance, in the work by Young et al [29] it is assumed that user goals don't change very much, that speech acts can be identified using intuition, and training data can be produced by "simulated users." Such work certainly results in interesting machine learning challenges, but is only peripherally relevant to better algorithms for human-machine conversation. Rather than modelling language itself as represented by a corpus, our approach models the language production process. This has been criticised by Wilks for only doing "one half of dialogue" (the non user's half) [30] but this is indeed how much of linguistics views the process. Our approach is theoretically well-founded, not in mathematics, but in linguistics. Rather than reasoning with uncertainty – a task for which statistical approaches are eminently suited – we find that humans reason about the uncertainties; rather than betting on the horses, human conversational tactics look more like the bookmaker and play the odds. Human language as used is full of contingency planning and mitigation strategies. The way to better machine dialogue systems is to find and copy those strategies.

The challenge therefore, is to capture understanding from the conversational dialogue, in form and detail suitable for a BDI plan library:

- the way an expert would fly a UAV in the context of the chosen scenario, and
- the way a human would explain the resultant plans-in-use to a human overseer.

With a system in place – either Wizard of Oz [31] or automated – we can record interactions and use those recordings to develop the system. Techniques for doing this have been developed as described, and related techniques have been developed for dialogue over the years including Conversation Analysis [25], Applied Cognitive Task Analysis [18] and, hopefully, our new Narrative approach [32].

5 Conclusions

There are many challenges with the scenario that we have described, but the opportunities offered by successfully deploying such a system are huge. As well as providing more effective search and rescue capabilities, similar systems can be used for exploration in difficult terrain, monitoring such things as traffic and large sporting events and if suitable data is available, large buildings and other structures that have met with some disaster. An example of this would be the rescue mission launched when the cruise ship Costa Concordia capsized. Search parties not only had the problems of exploring an unfamiliar enclosed environment, but it was at an unfamiliar angle and water levels had to be assessed and monitored. This was before the effects of damage had been taken into account. Small flying autonomous robots could rapidly explore the area and monitor human searchers, maintaining contact with the controller and warning them of potential dangers in good time, so that they can take evasive action, and keep themselves safe.

If a swarm of swimming robots with similar capabilities could also be deployed, the complete ship could be explored in a coordinated manner.

We believe that a limited conversational interface in conjunction with a Belief Desire and Intention architecture can provide an effective way to provide flexible operations, reduce work load, and minimise bandwidth requirements. Ultimately of course autonomous vehicles could look after themselves, providing only timely and relevant sensor data. With the state of the art, they need supervision and the proposed virtual environment provides a highly flexible human machine interface.

References

1. Mehlmann, G., Haring, M., Buhling, R., Wisner, M., Andre, E.: Multiple agent roles in an adaptive virtual classroom environment. In Allbeck, J., Badler, N., Bickmore, T., Pelachaud, C., Safonova, A., eds.: *Proceedings of the 10th International Conference on Intelligent Virtual Agents (IVA'10)*, Berlin, Heidelberg, Springer-Verlag 250–256
2. Gemrot, J., Brom, C., Plch, T.: A periphery of pogamut: from bots to agents and back again. In: *Agents for Games and Simulations II. Volume 6525 of LNCS.*, Springer (2011) 19–37
3. Richards, D., Taylor, M., Porte, J.: Allowing trainers to author their own learning content. In: *Proceedings of Workshop on Education and MAS (EduMAS09)*. (2009)
4. Saez-Pons, J., Alboul, L., Penders, J.: Experiments in cooperative human multi-robot navigation. In: *International Conference on Robotics and Automation (ICRA 2011)*, Shanghai, China, The IEEE Robotics and Automation Society (2011) 6512–6516
5. Saez-Pons, J., Alboul, L., Penders, J., Nomdedeu, L.: Multi-robot team formation control in the GUARDIANS project. *Industrial Robot: An International Journal* **37**(4) (2010) 372–383
6. Schank, R.C., Abelson, R.P.: *Scripts, Plans, Goals, and Understanding*. Lawrence Erlbaum Press (1977)
7. Allen, J.F., Schubert, L.K., Ferguson, G., Heeman, P., Hwang, C.H., Kato, T., Light, M., Martin, N.G., Miller, B.W., Poesio, M., Traum, D.R.: The TRAINS project: A case study in defining a conversational planning agent. *Journal of Experimental and Theoretical AI* **7**(7) (1995)
8. Wallis, P.: Robust normative systems: What happens when a normative system fails? In de Angeli, A., Brahmam, S., Wallis, P., eds.: *Abuse: the darker side of human-computer interaction (CHI 2005)*. (2005)
9. Wallis, P.: Believable conversational agents: Introducing the intention map. In Pelachaud, C., Andre, E., Kopp, S., Ruttkay, Z., eds.: *Creating Bonds with Humanoids (Proceedings of the Workshop at AAMAS'05)*. (2005)
10. <http://news.bbc.co.uk/1/hi/england/derbyshire/7765014.stm>
11. <http://www.bbc.co.uk/news/uk-england-derbyshire-11081145>
12. <http://www.bbc.co.uk/news/uk-england-derbyshire-12648565>
13. <http://www.bbc.co.uk/news/uk-england-stoke-staffordshire-13036916>
14. <http://www.bbc.co.uk/news/uk-england-south-yorkshire-17028351>
15. Penders, J., Alboul, L., Witkowski, U., Naghsh, A., Saez-Pons, J., Herbrechtsmeier, S., Habbal, M.E.: A robot swarm assisting a human fire fighter. *Journal of Advanced Robotics* **25**(1-2) (2011) 93–117
16. Payr, S., Wallis, P.: Socially situated affective systems. In Petta, P., Pelachaud, C., Cowie, R., eds.: *Emotion-Oriented Systems: The Humaine Handbook*. Springer (2011) 497–516
17. Tomasello, M.: *Origins of Human Communication*. The MIT Press, Cambridge, Massachusetts (2008)

18. Militello, L.G., Hutton, R.J.B.: Applied Cognitive Task Analysis (ACTA): A practitioner's toolkit for understanding cognitive task demands. *Ergonomics* **41**(11) (1998) 1618–1641
19. Norling, E.: What should the agent know?: The challenge of capturing human knowledge. In: Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems - Volume 3, Richland, SC, International Foundation for Autonomous Agents and Multiagent Systems (2008) 1225–1228
20. Rao, A.S., Georgeff, M.P.: BDI agents: From theory to practice. In Lesser, V., ed.: Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95), MIT Press (1995) 312–319
21. Heinze, C., Smith, B., Cross, M.: Thinking quickly: Agents for modeling air warfare. In: Proceedings of the Eighth Australian Joint Conference on Artificial Intelligence, Brisbane, Australia (1998) 47–58
22. Heinze, C., Goss, S., Pearce, A.: Plan recognition in military simulation: Incorporating machine learning with intelligent agents. In: Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence, Workshop on Team Behaviour and Plan Recognition, Stockholm, Sweden (1999) 53–63
23. Norling, E.: Modelling Human Behaviour with BDI Agents. PhD thesis, University of Melbourne, Australia (June 2009)
24. Alboul, L., Chliveros, J.: A system for reconstruction from point clouds in 3D: simplification and mesh representation. In: ICARCV 2010, 11th International Conference on Control, Automation, Robotics and Vision, Singapore, IEEE (2010) 2301–2306
25. Wallis, P., Mitchard, H., Das, J., O'Dea, D.: Dialogue modeling for a conversational agent. In Stumptner, M., Corbett, D., Brooks, M., eds.: AI 2001: Advances in Artificial Intelligence, Adelaide, Australia, Springer (December 2001) 532–544
26. Payr, S., Wallis, P., Cunningham, S., Hawley, M.: Research on social engagement with a rabbit user interface. In M., T., de Ruyter B., J., S., A., M., C., B., N., S., T., M., eds.: Roots for the Future of Ambient Intelligence. Adjunct Proceedings, 3rd European Conference on Ambient Intelligence (AmI09), ICT&S Center, Salzburg, Austria (2009)
27. Williams, J.D., Young, S.: Partially observable Markov decision processes for spoken dialog systems. *Computer Speech & Language* **21**(2) (2007)
28. Lemon, O.: Learning what to say and how to say it: joint optimization of spoken dialogue management and natural language generation. *Computer Speech and Language* **25**(2) (2011)
29. Young, S., Gasic, M., Keizer, S., Mairesse, F., Schatzmann, J., Thomson, B., Yu, K.: The hidden information state model: A practical framework for pomdp-based spoken dialogue management. *Computer Speech & Language* **24**(2) (2010) 150–174
30. Wilks, Y., Catizone, R., Worgan, S., Dingli, A., Moore, R., Field, D., Cheng, W.: A prototype for a conversational companion for reminiscing about images. *Computer Speech & Language* **25**(2) (2011) 140–157
31. Janarthanam, S., Lemon, O.: A wizard-of-Oz environment to study referring expression generation in a situated spoken dialogue task. In: Proceedings of the 12th European Workshop on Natural Language Generation. ENLG '09, Stroudsburg, PA, USA, Association for Computational Linguistics (2009) 94–97
32. Wallis, P.: From data to design. *Applied Artificial Intelligence* **25**(6) (2011) 530–548