

Outbreak Agent: Intelligent Wearable Technology for Hazardous Environments

Erika Rogers, Robin R. Murphy and Chris Thompson

ABSTRACT

The topic of this paper is a project to design an intelligent agent which supports humans in the domain of hazardous site investigation. This software agent will be integrated into wearable technology which is embedded in the protective clothing worn by the humans. An overview of some of the design issues is presented, with a particular emphasis on the visual problem solving component.

1. INTRODUCTION

Human exploration and/or intervention in environments which pose dangerous and life-threatening conditions will continue to be a necessity until current technology is sufficiently developed to rely on automated methods. In such diverse situations such as virulent disease outbreaks, nuclear and chemically toxic spills, or outer space exploration, the human's superior perceptual, manipulative and decision-making skills are employed with the aid of protective clothing, which acts as a buffer from the external noxious environment. These suits typically carry their own life-support systems to provide cooling and oxygen supply, and allow voice communication capabilities, but otherwise, their cumbersome nature poses challenges for the humans wearing them in trying to carry out delicate and/or difficult tasks.

The addition of features such as intelligent agency, multiple sensors and augmented reality through wearable computer technology, introduces the concept of *smart*

suits, which can not only enhance the safety and effectiveness of the users, but which may lead to a whole new methodology for conducting field work. In turn, this particular domain promises to provide a rich source of knowledge about how to characterize symbiotic human-agent relationships: where the human benefits through the augmentation of perception, problem-solving and knowledge navigation capabilities, while the software agent becomes situated through the human host, and is able to learn from interaction with the human in the task environment. Basic research is needed to investigate human and agent performance and learning in these tasks through cognitive studies, intelligent agent architectures, machine learning paradigms, innovative interface design, and hardware/software integration. A subset of these issues is addressed by the two main goals of the Outbreak Agent project:

1. Design a Performance Support Agent which demonstrates three major types of cognitively-based intelligence: visual problem solving assistance, reciprocal learning, and high level speech interface.
2. Integrate this software agent into state-of-the-art wearable technology hardware, which is embedded in a smart suit to be worn by humans in hazardous environments.

Some of the design issues are now presented, together with an overview of some of the system components. These are discussed in the context of two possible ap-

plication scenarios.

Scenario 1: Hazardous Material Inspections for Search and Rescue. Many natural disasters, such as earthquakes and flooding, have the potential to trap humans and release hazardous materials; consider earthquake damage at a pharmaceutical plant or a hospital. As with Oklahoma City, rescue workers must perform a structural inspection to see if it is safe for rescuers to enter, and must isolate potential hazards (i.e., natural gas leaks, industrial gas leaks, ruptured chemical drums, etc.). Often there are limited emergency frequencies and standard procedures to restrict radio communications during SAR activities, so that, under current conditions, the rescue worker can be unaided for much of the crucial activity.

Scenario 2: Epidemic Investigation and Disease Surveillance. With various epidemic outbreaks, workers are sent into an area to collect both data and specimens which will be used to determine what type of disease it is, what the medium of transmission is (e.g., is it food-borne, airborne, etc.) and what relationship it may have to other similar outbreaks. For much of this activity, protective clothing must be worn, not only to protect the investigators, but also to avoid contamination of the samples. Unfortunately, in remote areas, communication with the home base is often difficult, and there can be a time lag while the specimens and other data are transported or transmitted to the home base (e.g., CDC), analyzed, further tests requested, etc. While these field workers may be highly trained, individual differences in techniques may affect the consistency of the tests and samples. In the case of disease surveillance, often local medical personnel are recruited, who may not have specialized training in the methodologies needed for standardized monitoring.

2. BACKGROUND

While the idea of human-computer symbiosis has been around for almost forty years [2], it is only recently that technology has reached a degree of sophistication that makes this a concrete possibility. The development of *augmented reality* environments which employ see-through displays and projectors allows the environment to be enriched through “virtual annotation”: the physical properties of objects in the environment may be overlaid with electronic enhancements such as outlines and instructions (e.g., [3]).

The particular coupling of wearable technology with augmented reality allows humans to take their information spaces with them, and to perform their information processing “anywhere, anytime.” The implications for

usage in task environments where people do not typically sit at desks to perform the most important parts of their jobs are overwhelming. Applications in areas such as medicine (e.g., [1]), manufacturing (e.g., [6]) and even clerical work (e.g., [11]) show the great promise of this approach to human-machine interaction. However, many of these systems rely on the complementary notion of “ubiquitous computing,” i.e., where objects in the environment are “tagged” electronically (e.g., with barcodes [5]). This allows the wearable system to recognize and track objects in the surrounding area, but it restricts usage to fairly static and stable environments.

This approach, however, is not feasible for technology which will be deployed in field investigations where the layout and characteristics of the environment are unknown and often unexpected. In the domain of hazardous site investigation, the application of hands-free wearable technology is particularly appealing. However, due to the uncertain nature of the environment and possible restrictions on external communications, it is important to have as much “thinking power” on board as possible. Therefore, the main emphasis of this work is on the development of the intelligent software which collectively comprises the Performance Support Agent. Some of the unique features of this intelligent assistant, however, are specifically due to its realization in wearable hardware which accompanies the human to the site.

3. PERFORMANCE SUPPORT AGENT

The three main features of the performance support agent are now described in greater detail, with particular emphasis on the visual problem solving component.

Visual Problem Solving Assistance

The design of this component is based primarily on previous research on visual interaction in the domains of diagnostic radiology [7] and tele-assistance for remote robot supervision [8]. In that work, a blackboard architecture was used to coordinate human-agent interaction through cognitively-based perceptual and problem solving components. However, the focus of these systems was primarily on diagnostic tasks, whereas the Performance Support Agent for the smart suit needs to have more extensive capabilities. These include: visual aids to better view the external environment, context-sensitive automatic image enhancement, diagrammatic renderings of decision paths and concept hierarchies, data collection and summarization aids, intelligent decision support/problem solving aids, and multimedia reference and training aids. An overview of the three main visual interaction components of the Performance Support Agent (PSA) is shown in Fig. 1.




Figure 1: Visual Interaction Components of PSA

The **Attention Director** is responsible for coordinating those modules which provide advanced visual aids and help to focus the human's attention on important visual features. Examples of these include:

- *Augmentation module.* This module provides the more traditional augmented reality type of support in the form of overlays to real world structures and objects and graphical enhancements of the environmental features. It is expected that this functionality will be most useful in an external structured environment such as described in Scenario 1. For example, it could project layouts of what the structure should look like, the human's relative position, and superimpose lines where walls should be.
- *Autonomous Vision module.* While the human is moving around in the environment, a head-mounted video camera will be capturing live data and feeding it to one or more computer vision modules for automated feature extraction and object recognition. Since the images will be time-stamped, this module will let the Attention Director alert the human when it has spotted something that may have been missed. It is expected that this capability will prove useful in the search and rescue scenario, where part of the task involves moving through the damaged area, seeking particular objects.
- *Training module.* This module is particularly important for the less experienced user. In fact, in some cases, such as disease surveillance, the personnel are often local medical workers, who do not have a strong background in this type of specialized data collection.

In Kenya, training candidates for yellow fever surveillance are selected from district hospitals, health centers and mission hospitals. Training focuses on case recognition,

record keeping, use of a Yellow Fever case investigation form, handling of serum samples and obtaining exposure information [9].

In this type of situation, an online intelligent training capability would help ensure both accuracy and consistency, regardless of the user's experience level, or a high turn-over rate.

- *Annotation module.* This is a knowledge-based tool for image annotation which links domain knowledge and features of the current case under consideration to a set of "live" visual overlays for the image. Each overlay may represent an assessment by a different human expert, and multiple overlays can be displayed if desired. The annotations are constructed from a collection of relevant concepts, tools such as ROI and additional comments areas, and may further be hyperlinked to more extensive background knowledge or to other similar cases. The physical annotation itself is anchored at a position in the image selected by the annotator, but is movable around that anchor, so that the underlying piece of the image can always be examined.

In these types of tasks, the humans are frequently required to gather a huge corpus of data and sometimes physical specimens. Task-specific information is therefore an important aspect of the PSA, and this knowledge is organized by the **Task Director** component. Example modules include:

- *Context module.* This where knowledge is presented about what is considered "normal" in the setting. For example, in the medical scenario, standard information such as regional characteristics, village organization, medical care organization, local contacts, etc. would be captured. In a search and rescue scenario, this module would present information about the original building layout and other properties of the site such as where hazardous chemicals are supposed to be located, etc.
- *Expectation module.* The function of this module is to take what is known so far about the emergency situation (e.g., an outbreak, or a bomb explosion), and generate expectations about the conditions to be encountered. These expectations are then used to produce the first round of data collection templates incorporating these characteristics. For example, in the process of conducting an epidemic investigation, the field personnel often start with a data-gathering approach: interview a few patients in great detail about circumstances surrounding the disease onset such as what they ate and drank. From this information, a standardized interview questionnaire is constructed with which to gather more extensive data.




Figure 2: Summary of epidemic occurrences

In the context of Scenario 1, on the other hand, it could present knowledge (either stored or being provided by the emergency response team leader) about the semantics of the building (Room A is likely to have people, Room B should have industrial gases, etc.)

- *Observation module(s)*. These modules initiate and monitor the main data collection activities. They may take several forms: for example, one might use a high level speech interface to fill in the templates generated from the expectation module; another may allow the user to author additional templates, questionnaires and comments documents; a third module might be responsible for interpreting readings from attached micro-sensors, (e.g., presence of trace amounts of Material A).
- *Graphical Data Summary module(s)*. Once some of the data has been collected, it can be summarized and graphically displayed, both for reporting purposes as well as decision making. An example would be the automatic generation of a graph such as is shown in Fig. 2. The advantage of this functionality is that preliminary results would be immediately available, which might help to focus the further gathering of data and/or specimens. Trends could also be quickly detected, and possibly averted, instead of having to wait for the results to be processed by the home base and then communicated back to the site.

The main AI component of the PSA is found in the **Hypothesis Manager**, which is in charge of looking at the data, comparing the expectations with the observations, invoking additional declarative and procedural knowledge from the knowledge base, and posting hy-




Figure 3: Example of test results for spotted fever

potheses about the situation. These hypotheses may involve helping the user classify viruses from images, interpreting additional test results as in Fig. 3 [10], or invoking a decision tree:

...a patient is considered to have a case of Yellow Fever if two of the following symptom complexes are present: jaundice, hemorrhage, signs of encephalitis or renal involvement, a temperature ≥ 38 degrees C [9].

While these components have been described individually, it is expected that they will interact with each other in a cooperative fashion. For example, the expectation and augmentation modules can work together to annotate the augmented display, and to post any warnings: “do not enter without precaution Y”, advisements: “sensors L, M, and N should be turned down to prevent overloading, sensor O should be turned off”, and procedures. The autonomous vision module, in cooperation with the augmentation and expectation modules, may attempt to find landmarks which confirm the wearer’s location relative to the map. It could also employ object recognition methods to find and highlight (via the augmentation module), candidate objects of interest such as spills, body parts, etc., that the wearer might miss due to fatigue or concentration on other tasks (e.g., walking through rubble).

Numerous cognitive studies are planned throughout the design and development stages of the project to determine how people currently function in this type of scenario, and how they might change what they do in the presence of an intelligent agent. Interesting issues are raised, not only in the interaction between the smart suit and its human host, but also in the nature of the collaboration among all the “augmented” team members. Some questions to be addressed include:

- What are the tasks and how are they currently performed?
- What is the background knowledge needed to carry out these tasks? What is the background knowledge used during the field session (i.e., do they typically have to wait till they get back into the laboratory)?
- What is the nature of the communication between the human team members?
- How do the human team members learn from the outbreak encounters? Are there formal de-briefing sessions? Informal chat sessions? Job performance evaluations?
- How will the outbreak agent change the types and nature of the tasks? Will the assistance provided outweigh the physical and cognitive load of using the agent?

Modeling of the domain knowledge will enable us to build prototype tools for studying how the users navigate through the knowledge, and how different types of interfaces may facilitate that navigation. This will allow effective strategies to be identified for the display of task-related concept hierarchies, as well as problem solving schemas which the system can later “teach” the human wearer.

Reciprocal Learning

In this project we plan to investigate and model a concept which we call *reciprocal learning*, where the human learns from the agent, the agent learns from the human, and, finally, the agents learn from each other. If the user is new to the system, or is asked to perform a rare task for which full training may not have been available, training will be provided by the agent in terms of how to use the system itself, as well as the most effective problem solving strategies the agent currently knows for the particular task. In the early phases of “getting to know one another,” scaffolded intelligent tutoring techniques will be explored to critique the user’s session, to keep track of errors and oversights when they can be identified, and to consider whether the user’s deviations may not actually have some merit. As the user becomes more experienced in the tasks, the tutorial mode fades in favor of a more “help on demand” mode together with a debriefing dialogue at the end of each session.

On the other hand, the smart suit is expected to learn individual user preferences from its wearer, as well as new patterns of problem solving activity and knowledge navigation. It can then adapt its interface and some of its intelligent behavior to that of the human user.

Finally, the distributed learning among the agents themselves can take two forms: 1) online learning - if

the human discovers an effective way of coping in a critical situation, that information can be broadcast to other human-suit pairs; and 2) offline learning - patterns of training, interaction and problem solving used by the humans and agents for the same problems are generalized into new default schemas.

Technical approaches to these learning issues include the use of reinforcement learning, a technique which has recently met with success in robotics and situated agent applications [4], case-based learning to capture key aspects of problem solving strategies for new tasks, and theory revision as a mechanism of transferring previously learned knowledge to a new situation without interference to the old tasks.

High Level Speech Interface

The importance of providing an easy-to-use interface which minimizes the need for the human to manipulate input devices has led us to include a high level speech interface as a significant aspect of the smart suit’s capabilities. While current wearable technology typically includes some speech recognition in the form of off-the-shelf systems, the particular nature of our scenarios provides a unique opportunity for the study of speech interface issues. Some of the verbal communicative tasks which are of interest include: retrieval of information (textual, auditory or visual) that is related to the work the user is doing, and may be stored locally or remotely; recording of findings made in the hazardous environment, through entry of text, numbers or image annotations; navigation between different kinds of tasks and displays, such as switching from recording to retrieving; asking for help on how to accomplish any task, and asking for help in interacting with the smart suit itself (similar to just-in-time training techniques used in the manufacturing domain); and giving explicit instructions to the smart suit.

Many forms of learning, both implicit and explicit, will be part of intelligent interaction between users and systems in the future. Explicit verbal instruction has a powerful role to play in teaching the smart suit about new concepts, relations between concepts, and new procedures to perform.

4. WEARABLE TECHNOLOGY

The Factory Automation Support Technology (FAST) [6] is a wearable, voice-activated computer system currently under development at Georgia Tech Research Institute, and has been developed for the manufacturing domain. It is equivalent to a 486 75Mhz desktop computer, and consists of the following components: a see-through display which allows the user to work while looking at text, drawings, and video that are pertinent to the

user's job; a miniature microphone/earphone which provides audio information to the user and accepts voice input from the user, and keeps the user's hands free for job-related tasks; a wireless communications link which sends and receives up-to-date information to and from a main computer system; a wearable computer which allows the user to enter and receive information wherever the user goes; and a battery pack to supply power for all the components.

This computer system enables employees to get information at the task site and, since their hands are not busy operating the computer, to continue to perform a task as they are receiving the information. Interaction with this system currently permits users to request information from a local storage device, download current information from a central files server via the wireless network, collect and store local data for later transfer to a central computer system, transmit live data from a worksite to a central facility via the wireless network, connect with a remote location transmitting live audio and video in both directions, determine a location through an attached GPS, and access worldwide information sources through an internet gateway.

While much of the prototype hardware has already been developed, it remains a challenge to adapt the components to the special requirements of the protective clothing. Ergonomic issues involve placing camera, speaker and microphone inside the head piece in a way that can easily accommodate the human user, and designing additional input devices that enable cursor movement in a spatially-oriented display (e.g., selecting areas on an image). Basic sound and video recording devices should also be incorporated and time-synchronized with the intelligent displays, so that an entire outbreak session may be captured for later review. This also serves as knowledge acquisition for future work on video indexing and recognition research.

5. CONCLUSIONS

Although some aspects of the scenarios presented above may seem to be somewhat futuristic, the reality is that the hardware technology to support much of what has been suggested is already available. The critical factor is the development and integration of the advanced software components which will make the Outbreak Agent a viable field partner for the humans who are risking their lives in these environments. It is expected that the outcome of this project will have both practical and theoretical implications. On the one hand, it should reduce some of the challenges faced by the investigators in the field, and assist in solving problems more rapidly and effectively. On the other hand, it will provide an interesting new platform for study of both human-agent and agent-agent interaction.

6. ACKNOWLEDGEMENTS

This research is currently supported by the Army Center of Excellence in Information Science at Clark Atlanta University, ARO Grant # DAAL03-92-G-0377, and NSF Grant IRI-9502289. The authors would also like to thank Dr. Colleen Crangle for her contributions.

7. REFERENCES

- [1] M. Bajura, H. Fuchs and R. Ohbuchi, "Merging virtual objects with the real world: Seeing ultrasound imagery within the patient". *Computer Graphics* 26 (2): 203-210.
- [2] J.C.R. Licklider, "Man-Computer Symbiosis", *IRE Transactions on Human Factors in Electronics*, vol. HFE-1, March 1960, pp. 4-11.
- [3] S. Feiner, B. MacIntyre and D. Seligmann, "Knowledge-Based Augmented Reality". *Communications of the ACM*, Vol. 36, No. 7, July 1993, pp. 52-62.
- [4] J. Lindner, R. Murphy, and E. Nitz, "Learning the expected utility of sensors and algorithms." In *IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems*, 1994, pp. 583-590.
- [5] K. Nagao and J. Rekimoto, "Agent Augmented Reality: A Software Agent Meets the Real World". *Proc. of the Second International Conference on Multiagent Systems*, Dec. 1996.
- [6] J.J. Ockerman, L.J. Najjar, J.C. Thompson, C.J. Treanor and F.D. Atkinson, "FAST: A research paradigm for educational performance support systems". In P. Carson and F. Makedon (Eds.), *Educational Multimedia and Hypermedia 1996*, Charlottesville, VA: Association for the Advancement of Computing in Education, pp. 545-550.
- [7] E. Rogers, "VIA-RAD: A blackboard-based system for diagnostic radiology". *Artificial Intelligence in Medicine* 7, 1995, pp. 343-360.
- [8] E. Rogers, R. R. Murphy and B. Ericson, "Agent-Based Expert Assistance for Visual Problem Solving". *Proc. of 1st International Conference on Autonomous Agents*, New York: ACM Press, 1997, 156-163.
- [9] E.J. Sanders, P. Borus, G. Adamba, G. Kuria, P.M. Tukei, and J.W. LeDuc, "Sentinel Surveillance for Yellow Fever in Kenya, 1993 to 1995". <http://www.cdc.gov/ncidod/EID/vol2no3/sanders.htm>, *Emerging Infectious Diseases*, Vol 2, No. 3, Jul-Sep 1996.
- [10] B.L. Smoak, J.B. McClain, J.F. Brundage, L. Broadhurst, D.J. Kelly, G.A. Dasch and R.N. Miller, "An Outbreak of Spotted Fever Rickettsiosis in U.S. Army Troops Deployed to Botswana". <http://www.cdc.gov/ncidod/EID/vol2no3/smoak.htm>, *Emerging Infectious Diseases*, Vol. 2, No. 3, Jul-Sep 1996.
- [11] P. Wellner, "Interacting with Paper on the DigitalDesk". *Communications of the ACM*, Vol. 36, No. 7, July 1993, pp. 86-96.