R-CAST-MED: Applying Intelligent Agents to Support Emergency Medical Decision Making Teams

Shizhuo Zhu[1], Joanna Abraham[1], Sharoda A. Paul[1], Madhu Reddy[1], John Yen[1], Mark Pfaff[1], and Christopher DeFlitch[2]

[1] College of Information Sciences and Technology, The Pennsylvania State University, Unversity Park, PA 16802, USA
{szhu, jabraham, spaul, mreddy, jyen, mpfaff}@ist.psu.edu,
[2] Department of Emergency Medicine, Penn State Milton S. Hershey Medical Center, Penn State College of Medicine, Hershey, PA 17033, USA
cdeflitch@psu.edu

Abstract. Decision-making is a crucial aspect of emergency response during mass casualty incidents (MCIs). MCIs require rapid decisions to be taken by geographically-dispersed teams in an environment characterized by insufficient information, ineffective collaboration and inadequate resources. Despite the increasing adoption of decision support systems in healthcare, there is limited evidence of their value in large-scale disasters. We conducted focus groups with emergency medical services and emergency department personnel who revealed that one of the main challenges in emergency response during MCIs is information management. Therefore, to alleviate the issues arising from ineffective information management, we propose R-CAST-MED, an intelligent agent architecture built on Recognition-Primed Decision-making (RPD) and Shared Mental Models (SMMs). A simulation of R-CAST-MED showed that this tool enabled efficient information management by identifying relevant information, inferring missing information and sharing information with other agents, which led to effective collaboration and coordination of tasks across teams.

Keywords: Intelligent Agent, Simulation, Mass Casualty Incidents, Decision Support Systems, Information Management

1 Introduction

Emergency medical decision-making is complex, especially during mass casualty incidents. In a *Mass Casualty Incident* (MCI), hospitals are required to deal with a large influx of patients with various levels of trauma in a short period of time [1]. The two key goals for healthcare providers during such MCIs are rapid evacuation of patients from the incident site and provision of critical medical care to a large number of patients [2]. These activities require a coordinated effort on the part of pre-hospital and hospital-based teams. Medical decisions taken by these patient care personnel during such crises, both individually and collectively, significantly impact the mortality rate of critically injured patients.

Pre-hospital services such as police, fire, EMS (Emergency Medical Services) and HAZMAT (Hazardous Materials) teams are responsible for ensuring that patients are stabilized and transported rapidly. They need to resolve such issues as how many patients need to be transported, how to transport the patients, and which facilities are best suited to handle the patients. The decisions on where to transport patients are usually based on the trauma level of the patients and on particular emergency department (ED) capabilities. The EDs receiving these patients are required to make decisions such as, how to triage the large number of incoming patients, whether to seek assistance from other departments, and when to alert other EDs or request additional resources. Though the pre-hospital and ED teams make different kinds of decisions, there is a *decision dependency* between these different teams.

We conducted focus groups with both EMS and ED teams at a major teaching hospital in order to gain insight into the difficulties with MCI decision-making. Participants were presented with the scenario of a train derailment involving hazardous materials and asked to describe how they would react to events and the prominent challenges they would face. Based on their responses, we discovered that timely access to relevant information is not only a major requirement but also a major challenge, for decision-making during a MCI. For instance, information required by ED team to make decisions, such as how many beds to prepare for incoming patients, depends on information available to and provided by EMS team, such as how many patients are en-route to that ED. This *information dependency* plays a key role in decision dependency.

Computer-based decision-support systems have been used for clinical and administrative purposes in a variety of settings, but they have rarely been applied to decision-making in MCIs. Existing clinical decision support systems (CDSSs) are primarily used to facilitate decisions regarding a *single patient* and by a *single team* of healthcare providers, and thereby limited in their ability to deal with MCIs where decisions are made about multiple patients by multiple inter-professional teams. Based on our fieldwork, we propose an agent-based emergency medical decision support system, R-CAST-MED, to help healthcare providers deal with the challenge of information management during MCIs.

The following section provides background on decision-making in MCIs and the use of decision support systems in healthcare. Section 3 and 4 describe the architecture of R-CAST-MED and the simulation of a particular MCI scenario. In section 5, we discuss the significance of our simulation and the design recommendations to better support decision making during MCIs. Finally, we conclude with some thoughts on role of agents in medical decision-making and future work in section 6.

2 Background

A mass casualty incident is any situation or event that places a significant demand on medical equipment and personnel [3]. Healthcare providers involved in dealing with patients of a MCI have to deal with a variety of challenges including organizational, logistical, and patient-care related [1]. Their response to these challenges will effect the mortality rate of critically injured patients. The post incident analyses of major

MCIs such as the World Trade Center attacks in 2001 [3], the London bombings in 2005 [2], etc. have performed careful assessment of the response to these MCIs. The analysis found that decisions made during such incidents played a crucial role in the outcome of the incidents. For instance, during the World Trade Center attacks, the decision on where to locate the emergency management command post was a mistake and decisions taken by emergency responders to transport all the initial patients to the three nearest hospitals overwhelmed those institutions [3]. These kinds of MCIs highlight the complexity of the decision-making process during MCIs because of the large number of people involved, time pressure, and the uncertainty of dealing with a new situation.

Decision Support Systems (DSSs) have been employed in healthcare to serve different purposes. For example, IDEAS for ICUs [5] makes use of case-based reasoning tools for disease diagnoses of patients. Another application of DSSs has been in evidence-based medicine where it helps in providing access to relevant patient data and knowledge base needed to make timely decisions. Despite their value in these healthcare areas, there has been limited research on DSS applications for crisis management [6]. Some decision support applications for emergency response during a crisis situation are "iRevive", a mobile pre-hospital database system that supports point-of-care electronic patient data capture that assists in triage decision making [7]; "Automated Triage Management (ATM)", a decision support model that assists healthcare practitioners to find patients' chief complaints [8]; and "Mobile Emergency Triage (MET)", a DSS model designed for pediatric population [9].

Though these DSSs accelerate the clinical diagnosis process during MCIs, they do not explicitly support dependencies in work, such as filtering and sharing appropriate information among multiple professional teams.

3 R-CAST-MED

3.1 Focus Groups

To better understand the challenges associated with MCI decision-making as well as to examine ways to support and improve the same, we conducted 7 focus groups with EMS and ED personnel associated with a 500-bed teaching hospital. We presented participants with the following scenario of a train derailment incident.

> ***Scenario:*** *A 76-car Norfolk Southern freight train carrying hazardous materials derailed in Derry Township, Dauphin County. The track where the derailment occurred runs parallel to E. Hershey Park Drive and is close to the golf course of the Country Club of Hershey. Patients of this derailment are being brought into the ED while the ED is operating at capacity.*

The 21 participants included air and ground EMS, attending and resident physicians, and communication center personnel. We presented participants with the scenario, and asked them questions regarding their decision making process during a MCI.

We discovered that the presence of geographically distributed teams of pre-hospital and hospital personnel with varying goals, training levels, priorities, and information requirements increased the complexity of the situation. We also found that the decisions made by one team depended on the decisions made by other teams, This decision dependency arose primarily out of an information dependency; in order to make decisions each team needed information that could only be provided by the other team. In addition, the effective collaboration among teams was highly dependent on the relevant information shared among them. Decisions made by these teams during MCIs required up-to-date, accurate and relevant information to be exchanged between them. When presented with our train derailment scenario, some questions asked by ED physicians included "How many patients are involved?", "What is the acuity level of patients coming to the ED?" This incoming information would help ED team make decisions on how many ED beds and trauma bays to prepare, whether to set up decontamination tents, etc. The primary source of this information for them would be the communication center of hospital which would in receive this information from the on-site first responders. Thus, the information received from first responders at the incident site would enable decision-making in the ED.

The complex and dynamic nature of a MCI necessitates the need for a decision support system that is user-friendly with the flexibility to choose how information is sent, received, filtered and shared, depending on the context of crisis environment. To address some of the challenges we identified in the focus groups, we extended R-CAST (RPD-enabled Collaborative Agents Simulating Teamwork) [10] and developed R-CAST-MED (R-CAST for Medicine) to support healthcare providers in decision making tasks by filtering, proactively gathering, providing and sharing relevant information.

3.2 R-CAST

Cognitive Foundations. R-CAST is a collaborative agent architecture built on *Recognition-Primed Decision-making (RPD)* model and *Shared Mental Models (SMMs)*. *RPD* model [11] describes how experienced people make decisions in dynamic environments. This model argues that in complex situations human experts usually make a satisficing rather than an optimal decision choice based on the recognition of similarities between the current situation and previous experiences. *SMM* is a hypothetical cognitive construct that refers to a common understanding among team members regarding their objectives, roles, etc. SMM attempts to explain many of the human behaviors in high performance teams [12].

R-CAST Agent Architecture. R-CAST is a RPD-enabled collaborative agent architecture extended from CAST (Collaborative Agents Simulating Teamwork) [10]. From a software engineering perspective, R-CAST is a component-based configurable agent architecture, i.e. each agent is configured by enabling/disabling components depending on the particular application. This adaptive feature allows R-CAST to be well-suited for the medical domain.

Fig. 1 depicts the basic architecture of a R-CAST agent. The knowledge base manager, information manager, communication manager, RPD-based decision making manager, and process manager are the key components. The knowledge base, experience base, and plan library are the repositories that contain inferential knowledge, experiential knowledge, and procedural knowledge respectively. The knowledge base defines *fact types* that the agent understands, *rules* that the agent uses to infer new information, and primary *facts* that the agent has already possessed. The experience base comprises of tree-like *experience spaces*, where every single experience encapsulates *cues*, *expectancies*, *goals*, and *course of actions (COA)*. The plan library specifies how the agent executes the COA in the form of *plans* and *operators*. The *domain adapter* is the interface between an agent and its surrounding environment, which specifies domain-dependent functions and capabilities.

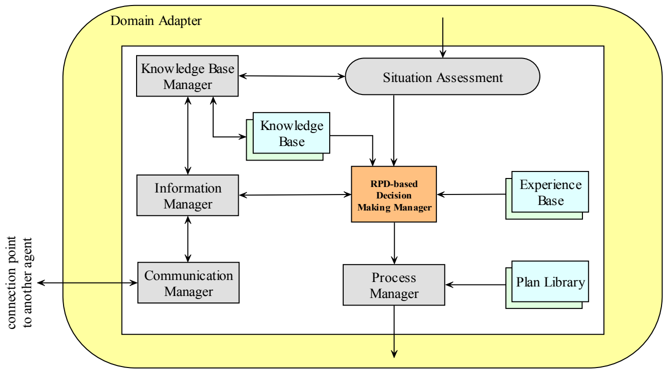

Fig. 1. R-CAST Agent Components

In general, the R-CAST agent updates the knowledge base with newly obtained information through *knowledge base manager* by constant observation and assessment of the situation. Meanwhile, feature matching is performed by comparing the current situation with the cues of existing experiences, by *RPD-based decision making manager*. If one satisficing experience is matched, then its corresponding COA is captured and executed by the *process manager*. If there is no match, then the *information manager* identifies the missing information requirements, and then sends an information request message to the *communication manager*. Upon receiving an information request message, the communication manager will communicate with other agents which are potential information sources.

3.3 R-CAST-MED: Adapting R-CAST to Emergency Medical Domain

The cognitive RPD model which describes the decision making process of an experienced person in emergent situations allows R-CAST to be applied in the emergency medical domain. R-CAST-MED utilizes and formalizes the information dependency feature of MCIs to support better decision making. This information dependency feature should be understood within certain *context*, which is formalized in R-CAST-MED in three forms: *inferential context*, *experiential context*, and *procedural context*.

The inferential context is represented in the knowledge base in two ways. First, the knowledge base predefines the type and format of information that is highly contextual which the agent can understand. This feature is considered as *contextual information relevance*. Second, the knowledge base defines the causal inference relationship among information, which is considered as an *inferential information dependency*.

The experiential context is represented in the experience base and is composed of the marking configurations of the active decision spaces, one for each space. The contextual content relies on the representation of cues, expectancies, COA, and anomalies, etc.

The procedural context reflects the relationships among processes, including preconditions, effects, and execution orchestration. The execution of an action depends on certain preceding conditions, and causes new effect to the situations. The procedural context is specified in the plan library.

Based on these contexts, R-CAST-MED utilizes information dependency feature to support effective information management. The information overload issue is alleviated by filtering irrelevant information. Requirements for missing information are identified by inferring lower level information from higher level information, which makes information seeking an objective-oriented process. Also, the information dependency across agents requires appropriate information to be distributed and shared. Therefore, R-CAST-MED allows healthcare teams including pre-hospital and hospital services to quickly process and fuse information from multiple sources in crisis management.

4 Simulation

The goal of this simulation was to examine how information was appropriately filtered, sought, and shared among agents, and how decision recommendations were made by agents depending on this information. We primarily focused on agents' abilities to interact with other agents. We used the scenario provided in section 3.1 as an input to our simulation. It was performed on a GUI (Graphical User Interface) platform adapted from NeoCITIES [13].

Based on the scenario, we built four agents corresponding to four teams (Fig. 2): 911 county communication center (911CCC), hospital communication center (HCC), EMS, and ED, in addition to four GUIs, one display for each agent. The R-CAST-MED agents were created by configuring their knowledge bases, experience bases, and plan libraries based on the scenario data. We employed a server to periodically generate event reports based on a predefined scenario text file. The server continuously sent various types of information, such as event location and number of patients to relevant agents (e.g., teams). For instance, the information that a patient is calling 911 for help would, in the simulation, be sent to the 911CCC agent. Therefore, based on the incident information received from the server or from other agents, the responsible agent made decision recommendations which were later displayed on the GUI.

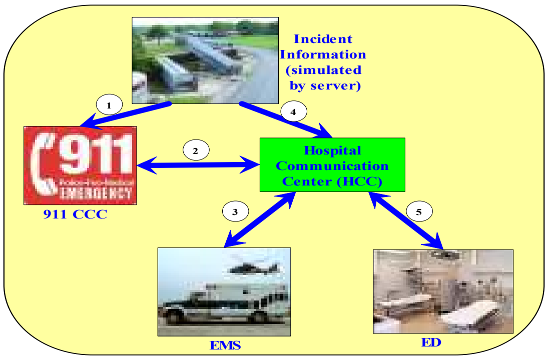

Fig. 2. Information Flow among Teams during a Train Derailment Scenario

The agents share some common knowledge, but differ in other kinds of knowledge specific to their respective context. To carry out certain types of tasks, an agent is required to know who would be the potential information source. For instance, EMS agent sought information regarding available ED resources from HCC agent, indicated by the following representation:

```
(FactType ED_resource(?type ?amount)
        (template "ED has ?amount much ?type type resource")
        (source
                (HCC plan_inform)
        )
)
```

Similarly, the HCC agent sought information from the ED agent. A chain of information seeking was thus created to capture the information dependency across different agents. The required information was delivered back to the requesting agent as soon as one of the requested agents in the chain had obtained it.

Agent recommendations were displayed on the GUI (Fig. 3), which is composed of four main panels: a map that locates the incident (upper right); a chat box for domain experts to exchange information (bottom right); an event tracker panel that provides the event description (bottom left); and an agent alert that displays agent decision recommendations (upper left). Features such as information seeking and sharing between agents is not depicted on the GUI.

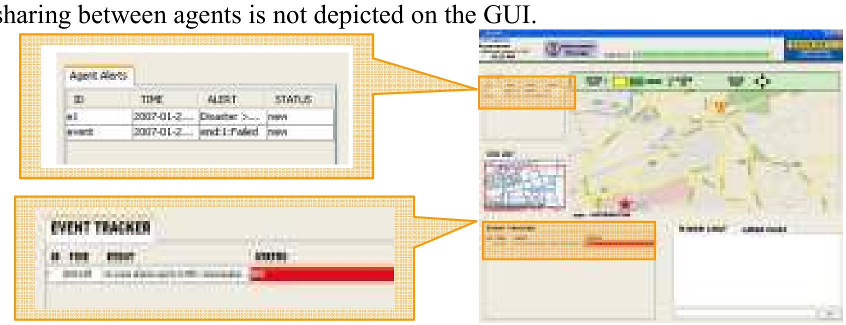

Fig. 3. Graphic User Interface for Displaying Disaster Scenario

As soon as the agent received an event report, the corresponding information was displayed on the event tracker panel. The agent recommendation was displayed on the agent alert panel. Fig. 3 shows an example where a MCI event with information on the number and severity of patients is reported and the agent recommends the ED agent "to activate the disaster plan" because of the high number of potential patients.

Our simulation demonstrated that agents can make decisions by effectively sharing and managing information. The results of the simulation showed that relevant and accurate information was exchanged between agents, and the appropriate decision recommendations were made. These decision recommendations were consistent with the data provided by the focus groups on what decisions they would make. We believe that this simulation highlights the potential for the R-CAST-MED agents to provide support for multiple teams to effectively collaborate and share appropriate and relevant crisis-related information without placing excess cognitive and affective constraints on the decision maker. For instance, in the simulation, the HCC agent responded to particular cues and delivered the information about the event location to the EMS agent without overloading the EMS agent with other irrelevant and extraneous information.

5 Discussion

5.1 Supporting Decision Dependency and Multi-Team Decision Making

For intelligent agent systems to play a useful role during a MCI, they must be able to facilitate and support decision dependencies and multi-team decision making. In the following paragraphs, we use real-world examples to illustrate how R-CAST-MED supports these two key features.

The EMS after arriving at the scene, assesses the situation at the incident site to decide "whether to transport patients to ED" and "how many patients that the ED can accommodate". In order to accomplish this goal, the *decision making* component of the EMS agent compares the assessment of the current situation (e.g. number of patients) with the experiences in the experience base (e.g. how many patients should be transported to the ED, whether the patients need immediate trauma care). This component chooses one of the two paths: 1. if there is a match, the decision choice will be made and its corresponding COA would be selected from the plan library and executed by the process manager; 2. if there is no match, it will request the information manager for missing information (e.g. ED resource availability); in cases when the information manager cannot find such information in its local knowledge base, it requests the communication manager to contact another agent (which is affiliated to another team) for this missing information critical for decision making.

As illustrated in the above example, decisions are interrelated because a decision regarding the transportation of patients to ED is dependent on the information provided by the ED agent (through communication center agent) to the communication manager of the EMS agent. Decision dependency feature is reflected in R-CAST-MED in several forms including contextual information dependency (derived from situation); inferential information dependency (based on rules built in knowledge base); and team-across information dependency (arises from communication across teams).

The second distinguishing feature of R-CAST-MED is its ability to support multiple team decision making. R-CAST-MED can be used by teams composed of different professionals with varying skill levels that provide integrated care during a dynamic situation resulting in an influx of multiple patients. For example, The EMS agent furnishes the ED agent with details about patients' medical history, vitals and also, performs initial triage at the incident site prior to transport. Upon receiving this information from EMS agent, the ED agent can make necessary arrangements for patients that can be directly assigned to beds without repeating the triage process. The coordination support among multiple teams provided by R-CAST-MED leads to better quality of patient care given the rapid nature of the situation.

5.2 Designing Decision Support Systems to Support MCI

The chaotic and dynamic nature of MCIs causes inadequate access to relevant information, ineffective inter-team collaboration, isolated and redundant activities, communication breakdowns, and other affective and cognitive overload. To be effective in these environments, we need to design decision support systems (DSSs) that have (1) better contextualization features and (2) more proactive and rapid learning capabilities.

First, context is gaining increased attention as we are moving towards a more dynamic and integrated health system. Understanding the context of the information need based on the complexity of the situation is an important requirement for DSSs. For instance, many DSSs such as R-CAST-MED incorporate some contextual features. However, they still lack robust temporal and spatial contextual features that allow them to adapt to varying dynamic situations. Therefore, we must develop DSS that incorporate context in a meaningful way.

Second, by improving the learning ability of DSSs, we can support decision making in varying environments. Supporting the learning feature helps in identifying hidden associations in both explicit and implicit information that could be temporally and spatially distributed. This learning requirement necessitates DSSs to proactively synthesize new knowledge based on their ability to retain and recollect from past experiences. To support learning of DSSs, we need to understand human learning processes. In addition, learning algorithms such as Bayesian learning and case-based learning should be examined in order to verify its applicability in dynamic situations.

6 Conclusion

The decision making process during MCIs is complex in nature. There are multiple factors that influence the decisions made by emergency responders including the dynamic nature of the incident, the need to access relevant information rapidly, sharing of accurate information, resource constraints, and coordination among teams. In this paper, we investigated a prominent challenge of a MCI that deals with effective information management. To address this challenge, we developed R-CAST-MED, a decision support system that achieved effective agent-agent interaction. Based on our simulation of R-CAST-MED, we confirmed that it helps in supporting effective information management therefore leading to better coordination of care.

Although our simulation highlighted the agent-agent interaction in R-CAST-MED, we did not evaluate the human-agent interaction. In our future research, we plan on incorporating human decision makers into our evaluation to verify whether the system can assist humans in improving situation awareness and decision making effectiveness.

Acknowledgments

The authors would like to thank Victor Pilewski for his support during data collection process.

References

1. Hirshberg, A., Holcomb, J.B., Mattox, K.L.: Hospital Trauma Care in Multiple-Casualty Incidents: A Critical View Annals of Emergency Medicine 37(6) (2001) 647-652.
2. Aylwin, C.J., Konig, T.C., Brennan, N.W., Shirley, P. J., Davies, G., Walsh, M.S., Brohi, K.: Reduction in Critical Mortality in Mass Casualty Incidents: Analysis of Triage, Surge, and Resource Use After the London Bombings on July 7, 2005. Lancet, 368(9554) (2006) 2219-2225.
3. Asaeda, G.: The Day that the START Triage System Came to a STOP: Observations from the World Trade Center Disaster. Academic Emergency Medicine, 9(3) (2002) 255-256.
4. Oster, N., Nierenberg, R., Menlove, S., Chason, K., Pruden, J.: Reflections: September 11, 2001 – What We Learned. Academic Emergency Medicine, 9(3) (2002) 216.
5. Frize, M., Trigg, H.C.E., Solven, F.G., Stevenson, M., Nickerson, B.G.: Decision-Support Systems Designed for Critical Care. AMIA Fall Meeting. (1997)
6. Stephenson R., Anderson P.S.: Disasters and the information technology revolution. Disasters 21 (1997) 305–334.
7. Gaynor, M., Seltzer, M., Moulton, S., Freedman, J.: A Dynamic, Data-Driven, Decision Support System for Emergency Medical Services. ICCS 2005, LNCS3515 (2005) 703–711.
8. Guterman, J.J., Mankovich, N.J., Hiller, J.: Assessing the effectiveness of a computer-based decision support system for emergency department triage. IEEE, (1993).
9. Michalowski, W., Rubin, S., Slowinski, R., Wilk, S.: Mobile Clinical Support System for Pediatric Emergencies. Decision Support Systems, 36 (2003) 161-176.
10. Yen, J., Yin, J., Ioerger, T.R., Miller, M., Xu, D., Volz, R.A.: CAST: Collaborative Agents for Simulating Teamwork, in Proceedings of the 7th International Joint Conference on Artificial Intelligence (2001) 1135-1142.
11. Klein, G. A.: Recognition-primed decisions. Advances in man-machine systems research. W. B. Rouse. Greenwich, CT, JAI Press. 5 (1989) 47-92
12. Cannon-Bowers, J.A., Salas, E. and Converse, S.A.: Cognitive psychology and team training: Training shared mental models and complex systems. Human Factors Society Bulletin, Vol. 33 (1990) 1-4
13. McNeese, M.D., Bains, P., Brewer, I., Brown, C., Connors, E. S., Jefferson, T., et al.: The NeoCITIES simulation: Understanding the design and experimental methodology used to develop a team emergency management simulation. Proceedings of the 49th Annual Meeting of the Human Factors and Ergonomics Society (2005) 591-594.