Evaluating Plans and Human Response Using a Normative Multi-Agent System

Lauren Thévin, Julie Dugdale, Olivier Boissier, Catherine Garbay

Evaluating Plans and Human Response Using a Normative Multi-Agent System

Lauren Thévin
LIG CNRS, UGA (University Grenoble Alpes), France
lauren.thevin@imag.fr

Julie Dugdale
LIG CNRS, UGA (University Grenoble Alpes), France
University of Agder, Norway
julie.dugdale@imag.fr

Olivier Boissier
Laboratoire Hubert Curien CNRS, Institut Henri Fayol, France
olivier.boissier@emse.fr

Catherine Garbay
LIG CNRS, UGA (University Grenoble Alpes), France
catherine.garbay@imag.fr

ABSTRACT

We describe a socio-technical system designed to train different organisations in emergency management during the preparedness phase. In this system, both humans and the system work together in building organisational context awareness. The system uses a normative multi-agent system and a tangible table for user interaction. Real time virtual feedbacks are used to draw users' attention to the validity of their actions with respect to the emergency plan. Feedback allows the actors to be aware of the other factors that impact the validity of their action (actions of other actors, environment etc.) and to identify the behaviour that is expected of them by the other organisations. The system is still in its prototype phase. Its behaviour is illustrated by example scenarios, showing that it is possible to support collaboration amongst distant actors, in a way that only relevant information is shown.

Keywords

Crisis management, training, scenario, collaboration, tangible interaction.

INTRODUCTION

Training is essential in order to continually test and improve the response to crisis situations before they actually occur. Playing out crisis scenarios allows the various actors to act out the management plan, refine their response and improve communication and collaboration. We present an approach for training where the emergency plan and the responders behaviours and collaborations are assessed, in preparation phase. It is necessary to assess the plan since any errors that may contain will not always be detected at design time, but at run-time when the actors are playing out the live simulation. Conversely, it may very well happen that the actors face some unexpected situation resulting from the environment dynamics or actions of other actors and they do not follow the plan. We focus on coordinating a response orchestrated by several geographically remote actors. For example, a flood emergency that involves the coordination of diverse actors, such as the mayor, actors in the support cell, and actors in the logistics cell. There are several important issues to consider, such as: who should be informed about an event, action, or decision? What is the relevant information that must be transmitted? How can context awareness be assured? What kind of real-time feedback should be given to the actors? Our goal is to build shared organisational context awareness between remote actors to help them carry-out the emergency plan, highlighting any behavioural errors if necessary. Simply displaying the overall organisational context would lead to information overload for the actors. Instead, we choose the approach of highlighting the differences in understanding that the various actors may have about the organisation of the emergency. In essence this exposes the underlying issues that lead to behavioural and collaboration errors.

STATE OF THE ART

The goal of this work is to design and develop a collaborative support tool for crisis management. Specifically, we are interested in using a training tool to support the work of geographically distant actors, with a focus on their organisational context. The success of response operations is greatly affected by effective inter-organisational coordination (Perry, 1991). One of the main challenges is being able to account for the different perspectives among distant actors, including identifying conflicts and proposing possible solutions. Large-scale emergency response necessitates the collaboration among individuals belonging to different organisations and having different goals, backgrounds and procedures. This issue is closely related to organisational and communication protocols i.e. the norms followed by different stakeholder groups (De Brito, Thévin, Garbay, Boissier and Hübner, 2015) (Wijnqaards, Kempen, Smit and Nieuwenhuis, 2006).

Training tools for emergency managers are typically geared towards validating emergency plans and assisting collaboration amongst several organisations, and finally evaluating the outcome of strategic decisions. Validating emergency plans is typically done using walk-throughs in which case it is not easy to identify small problems with operational details. Otherwise plans may be validated during emergency drills and exercises (Perry and Lindell 2003); this can interrupt the flow of the exercise and disrupt training. These tools often rely heavily on simulation and some also allow collaborative training. For example, PANDORA uses a complex digital support environment for training commanders in strategic planning (Bacon and MacKinnon, 2011). ADMS is an emergency and disaster management training system designed to train incident commanders, first responders, and incident command teams in a real-time, interactive virtual reality environment (Louka and Balducelli, 2001). These, and other tools such as Hazmat Hotzone (Losh, 2007) and fire-fighter training simulator (Dugdale, Pavard and Pallamin, 2004) tend to focus on training individuals in only one organisation; training for collaboration across organisations is rarely considered. An exception is CRIMSON that uses virtual reality simulation to train crisis managers, field crews, etc. to plan and rehearse crisis missions (Balet, Duysens, Compdaer, Gobbetti and Scopigno, 2008). Likewise, INDIGO allows for cross-organisational exercises (Ahmad, Balet, Boin, Brivio, Ganovelli, Gobbetti and Schraap and 2012). However, these works focus on creating a common operational picture, via knowledge and information sharing; they do not offer any formal representation and handling of human emergency plans. Outside of the simulation domain, the Virtual Collaboration Environment (VCE) looks at distributed crisis response planning (Wickler, Potter, Tate and Hansberger 2011). Here subject-matter experts and crisis responders are asked to collaboratively create emergency response plans in order to better produce better solutions. Like VCE, COLLARIO allows virtual teams of geographically distributed experts to create and discuss emergency scenarios (Yao, Turoff and Hiltz, 2010). However, both VCE and COLLARIO focus on how collaboration can be supported, and not at resolving the underlying issues as to why a common operational picture (COP) cannot be formed.

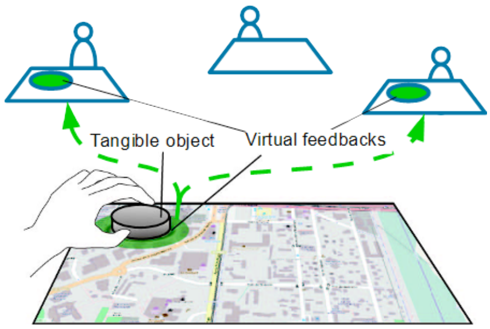

Figure 1. Users follow an emergency plan by interacting with tangible objects. The system analyses their actions in terms of their relevance to the plan and the quality of collaboration, and displays local or remote virtual feedback

Our goal is not to support the collaboration of people who are doing the same task in different geographical locations, nor to support the collaboration of people doing different tasks in the same location. Instead, we focus on supporting distributed collaboration where different organisations are jointly trying to achieve an overall task through an intricate coordination of their individual sub-tasks. This is achieved by allowing the actors to work with tools, such as maps, with which they are already familiar. We have designed a socio-technical system where both humans and the system work together in building organisational context awareness (figure 1). The training tool is in the form of a tangible table, called *Tangisense,* which displays a map of the emergency area. A variety of tangible objects may be used, some related to the geographical context of the map (e.g. a car), or not (e.g. the transmission of information). This form of interaction preserves the actors autonomy with respect to the plan they want to follow. Each tangible action performed by a given actor is observed on a step-by-step basis by the coordination agents, and analyzed with respect to the expected plan, considering the actor's role, mission and norms. In this way, any action, expected, permitted or not, is fully analysed. The normative system operates through posterior evaluation rather than a priori constraint. The table can be connected over a network and can display local and remote virtual feedback. These feedbacks are used (i) to ensure sharing of information over the network and (ii) to reflect the normative analysis performed by the agents, drawing the actor's attention to the validity of their actions.

APPROACH AND FORMALISM

Remote collaborations between actors from different organisations greatly increase the challenge of context awareness. Furthermore if actors have different interpretations of the organisational situation then a common understanding is less easily obtained. It is extremely useful to detect and display these differences in understanding between the actors since these represent conflicts and misunderstandings. Human actors typically do not interact specifically about a conflict, indeed they are often unaware that a conflict exists until some time later. In our system, the human actors are left free to undertake actions in order to respond to an emergency event. These actions are then analysed, in the scope of the plan and in relation to other users actions or the environment, to detect and display conflicts to the relevant actors. The users, now being aware of the conflicts, can work towards building a common interpretation of the organisation. Three cases must be considered to properly evaluate the actions of human actors: (i) the plan may be wrong, (ii) the plan is adequate but improperly followed, or (iii) there are some unexpected circumstances that lead the actor to deviate from the plan. The system uses a normative multi-agent system (MAS) to follow human activity and manage tangible and virtual interactions on the *Tangisense* table. A MAS is a computer program composed of autonomous interacting intelligent agents. The collective behaviour of these agents helps to manage a particular problem; in our case it is ensuring the good collaboration of geographically remote human actors. A normative multi-agent system uses the notion of norms to regulate agents' activity. A norm is a social or organisational rule. The normative dimension of the MAS allows us to represent rules and protocols of human crisis management activity and the users' interactions mediated by the system. Norms are used in three ways during the execution of a plan: to analyse and interpret users tangible interactions, to follow human activity according to the plan and the quality of the collaboration, and to display and manage feedbacks.

The approach is implemented using the JaCaMo multi-agent framework (Boissier, Bordini, Hübner, Ricci and Santi, 2013) that combines three separate technologies: Jason - for programming autonomous agents, Cartago – for programming environment artifacts, and an extended Moise – for programming multi-agent organisations. In JaCaMo formalization, a system S = <A,E,I,O>, where A refers to the agents, E to the artifacts (the representation of the tangible object), I to the interactions, and O to the organisations. A norm is an organisational rule that links a role to an action in a defined context. More formally, a norm ni is defined as ni = ⟨act-c, trigger-c, group, role, dm, object⟩, where act-c is the context of activation (e.g. suspicion of flood); trigger-c is the trigger event (e.g. a particular water level of the river); group refers to an organisation group (e.g. Command); role represents a role in the organisation (e.g. Leader); dm is the deontic attribute (e.g. obligation); and object is the expected action (e.g. evacuate). We define three functional spaces of interaction, adapted from the CLOVER groupware model (Laurillau et al. 2002): production, collaboration, and communication. The production space refers to a crisis intervention action and the tangible actions that a human actor can perform (e.g. evacuating an area by placing the tangible object "evacuate" on the table). The collaboration space analyses the action, assessing whether or not it is valid in terms of task and role assignment. Based on the interpretation of the analysis, the communication space deals with sharing relevant information, providing local and remote feedback to the appropriate actors. There is one production agent and one coordination agent per user, and one communication agent per table. The production agent is used to interpret activity associated with the tangible object. The coordination agent is used to follow and analyse the user's activities, and each agent follows one user. To check that the users follow the plan, the coordination agent uses organisational norms formalising the city's safeguard plan (PCS). For example a coordination agent will "evacuate", and check if it is a valid action of the plan by looking whether norms have been followed or violated. To check the presence of exceptions, that is actions that do not follow the plan due to unexpected circumstances, we currently enforce dynamic constraints on the action to be done in the plan, such as cardinality (how many people are allowed to execute the same action at a given time). The communication agent sends virtual feedback via the tables to actors' using the organisational rules to define the form of the feedback and the receivers. E.g. the communication agent receives an analysis of "valid evacuation action by the logistics cell". The communication agent will put a green circle under the "evacuate" tangible object that is on the logistic cell's table, and put an "evacuation" logo on the remote table of the Command Post (figure 1).

role	deontic	mission	time contraints
roleobssc roleobslc	permission	Mission declare magnitude observed noflood	
roleobssc roleobslc	permission	Mission declare magnitude observed flood 5	
roleobssc roleobslc	permission	Mission declare magnitude observed flood 10	
roleobssc roleobslc	permission	Mission declare magnitude observed flood 100	
roleobssc roleobslc	obligation	Mission declare flood magnitude	
rolechief	permission	Mission launch plan noflood	
rolechief	permission	Mission launch plan flood 5	
rolechief	permission	Mission launch plan flood 10	
rolechief	permission	Mission launch plan flood 100	
rolechief	obligation	Mission organize security	
rolemayor	obligation	Mission representative	
rolechief	obligation	Mission continue organize crisis	
rolerepresentative	obligation	Mission transmit crisis situation information	Every 15 min
rolefirefighter	obligation	Mission continue organize safeguard	

Figure 2. A: City safeguard normative specifications. According to the first 4 rows, the logistic cell and support cell observers are permitted to declare the observed magnitude of a flood. According to row 5, declaring the flood magnitude is obligatory.

EXPERIMENTAL PROTOCOL AND RESULTS

We have been working closely with the Major Risks Institute of Grenoble (IRMa) in order to understand emergency plans and to define our test scenarios. IRMa is a regional association concerned with training and providing advice to local decision-makers about their emergency plans and training exercises. More specifically, we use real exercise scenarios, a real PCS (city safeguard plan) and real experience feedbacks from both exercises and real crisis management situations. In a flood crisis management exercise, the city tests some part of its PCS. The management of the emergency is structured around one communal command cell (CCP), directed by the local mayor, and two field cells, a support cell (SC) and a logistics cell (LC), directed by adjuncts. We describe two scenarios that are modeled with the Moise normative specifications shown in figure 2. (Boissier *et al*., 2013) gives more information about the complete Moise organisational specification used. The missions referred to in the norms correspond to goals in the PCS: once the flood magnitude is declared, e.g. « goalflood10declared », the information is transmitted and the corresponding plan is launched (« launchedflood10plan »); security will then be achieved (« goalorganizedsecurity »). The mayor coordinates with the firefighters, e.g. by defining a representative (« goalcoordinationFF » and « goaldefinedrepresentative »). The firefighters have two roles (firefighters and city representative) and two goals, safeguard actions and regular transmission of information ("goalinformationtransmitted"). The first one shows checking exceptions resulting from an information conflict. The second one shows the violations resulting from inter-organisational collaboration issues. The interface of our system has not yet been fully developed therefore in the figures below we only show the results as they appear in code in our prototype.

Scenario 1: information conflict

In this scenario, two field cells observe the evolution of a flood using two different indicators. The SC is responsible for observing how high and fast the river level evolves; the LC is responsible for observing any overflow in a specific place and how fast this appears. They then qualify the flood as being centennial, decennial or quinquennial (i.e. as the flood was 5, 10 or 15 years ago), or no flood. The mayor must decide which (unique) plan to execute, following the differing information from the field cells. The system must be able to track the remote activities of the two field cells and detect whether they both followed the PCS. It should also recognise whether a conflict appears when it looks at the two different analyses. In the running scenario, due to differences in their analysis process, the LC qualifies the flood as quinquennial, while the SC qualifies it is as decennial, as shown by the upper red rectangles in the left of figure 3. This results in two concurrent incompatible goals for the mayor, as shown by the red ellipse in figure 3, that are correctly detected by the system.

Scenario 2: inter-organisational collaboration

In this scenario, the mayor is represented as playing two roles in two organisations: chief of the CCP and representative of the CCP for the firefighters. The firefighters expect the CCP to regularly send information reports and the mayor is expected to play this role. If this information is not sent in time, a conflict appears. Every organisation should be able to design and follow their plan independently, with the minimum amount of cross alignment. Hence the system should be able to model and test the consistency of the plans of two organisations; this is a common requirement in crisis management preparation. This running scenario shows that an actor may play two different roles in different independent organisations and be committed to different goals of which they are not fully aware. Figure 4 shows how the system correctly detects these goals and checks whether there are followed or not by the mayor, based on the specific norms attached to each organisation.

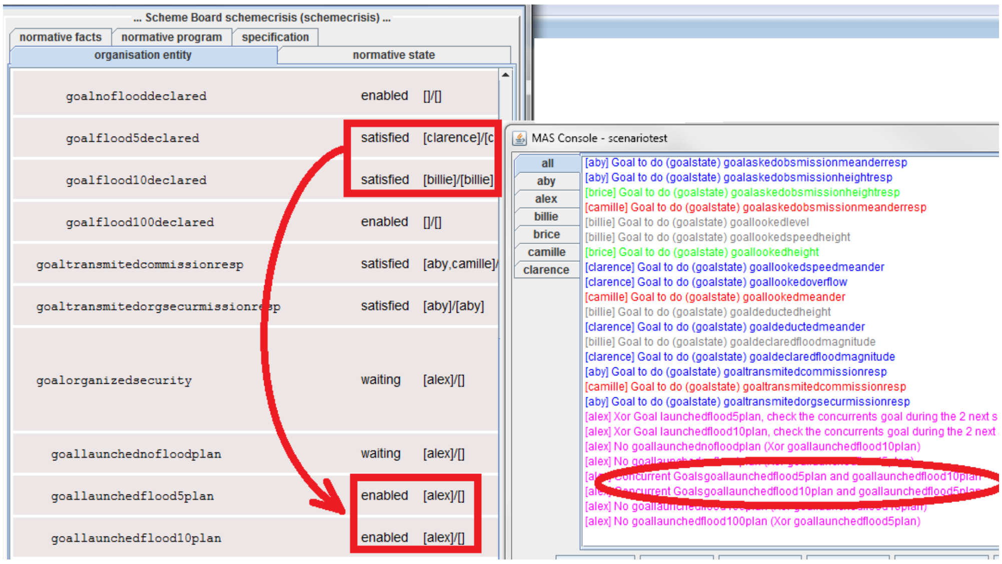

Figure 3. In the upper red rectangle, Clarence from the LC categorised the flood as quinquennial, whereas Billie from the SC declared it as decennial. Consequently, Mayor Alex's actions (bottom red rectangle) following the PCS are to declare two flood-plans that are consistent with the previous declarations. Declaring a flood plan is a valid action, but only one plan at a time can be launched, as stipulated by a dedicated constraint. The red ellipse shows the system detecting the conflict between the 2 flood plans, based on this constraint.

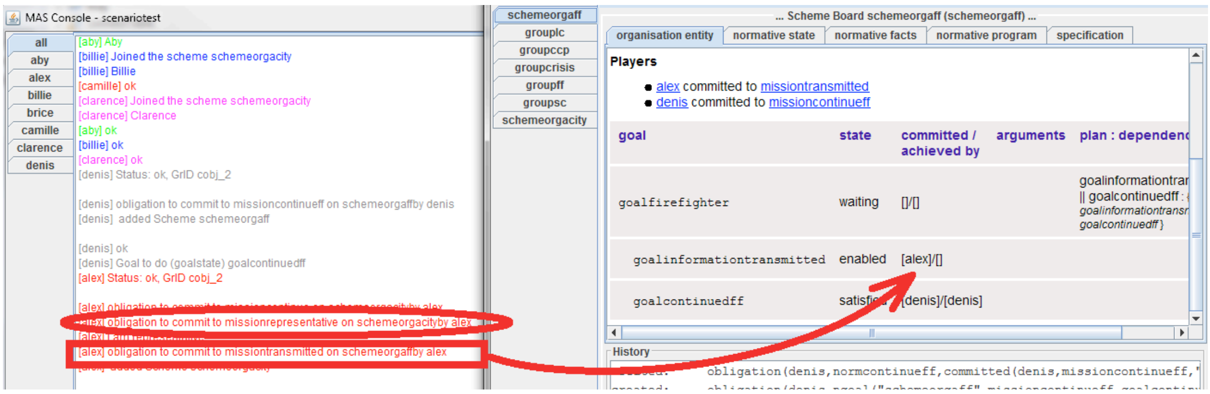

Figure 4. Alex, the Mayor of the city, is also a city representative to the firefighters. He commits to this mission (red ellipse). Accepting this role means that Alex also has responsibilities in the firefighters organisation (red rectangle). The red arrow indicates that it is correctly identified in the follow-up of the actor obligations.

CONCLUSION

We have shown some preliminary work toward the design of a training system able to support emergency management. The results showed how the system can follow and analyse human activities in responding to an emergency, at the plan level and at the collaboration quality level. The system can follow distributed but dependent activities, deal with shared resource usage, and handle the different points of view and plans of various organisations. It uses organisational context awareness to support the actors in a situated way. The actors can obtain feedback on their actions in real-time, check the validity of their actions with respect to their plan, be aware of the other factors that impact the validity of their action identifying the behaviour that is expected of them by other organisations. This preliminary work has demonstrated that it is possible to have a training system that can support collaboration amongst distant actors, in a way that only relevant information is shown. Future work includes enhancing both analysis of activity and the interactional functionalities. We could, in the coordination space, analyse in more depth the reason behind exceptions. E.g. in the first scenario, we could detect which of the previously linked actions led to the exception situation. This may also be a good basis for feedbacks, as long as the mayor decides to follow a plan that is not aligned with at least one of the local visions. Then, we will develop the interaction technical structure to compute tangible interaction and feedback delivery. Different levels of abstraction will be created, as long as situated activities needs meaningful aggregation of tangible objects, humans and feedbacks activity.

ACKNOWLEDGMENTS

The authors thank the members of IRMa (Major Risks Institut in Grenoble) for their collaboration and the Rhône-Alpes Region for their support for the ARC6 contract (ARC-13-009716-01).

REFERENCES

1. Ahmad, A., Balet, O., Boin, A., Brivio, P., Ganovelli, F., Gobbetti, E. and Schraap, M. (2012) Interactive Simulation Technology for Crisis Management and Training: The INDIGO Project, *Proceedings of the ISCRAM 2012*, Vancouver, CA.
2. Bacon, L, Wandall, G. and MacKinnon, L. (2011) The development of a rich multimedia training environment for crisis management: using emotional affect to enhance learning, *Proceedings of the ALT-C 2011*, Leeds, UK.
3. Balet, O., Duysens, J., Compdaer, J., Gobbetti, E. and Scopigno, R. (2008) The Crimson Project Simulating populations in massive urban environments. *Proceedings of the WCCM8 2008*, Venice, IT.
4. Boissier, O., Bordini, R. H., Hübner, J. F., Ricci, A. and Santi, A. (2013) Multi-agent oriented programming with JaCaMo, *Science of Computer Programming*, 78, 6, 747-761.
5. De Brito, M., Thevin, L., Garbay, C., Boissier, O. and Hübner, J. F. (2015) Situated Regulation on a Crisis Management Collaboration Platform, *Proceedings of the PAAMS 2015*, Salamanca, ES.
6. Dugdale, J., Pavard, B. and Pallamin, N. (2004) Emergency fire incident training in a virtual world. *Proceedings of the ISCRAM 2004*, Brussels, Belgium.
7. Laurillau, Y. and Nigay L. (2002) Clover architecture for groupware, *Proceedings of CSCW 2002*, New Orleans, Louisiana.
8. Losh, E. (2007) The birth of the virtual clinic : Games spaces in « the virtual praticum » and the Virtual Terrorism Response Academy, *Proceedings of the ISCRAM 2007*, Delft, NL.
9. Louka, M. and Balducelli, C. (2001) Virtual Reality Tools for Emergency Operation Support and Training, *Proceedings of the TIEMS 2001*, Oslo, NO.
10. Perry, R.W. (1991) Managing Disaster Response Operations. In T.E. Drabek and G. Hoetmer (eds.) Emergency Management: Principles and Practice for Local Government, International City Management Association, Washington.
11. Perry, R.W. and Lindell, M.K. (2003) Preparedness for emergency response: Guidelines for the emergency planning process, *Disasters*, 27, 4, 336-350.
12. Wickler, G., Potter, S., Tate, A. and Hansberger, J. (2011) The virtual collaboration environment: New media for crisis response. *Proceedings of 8th ISCRAM*, Lisbon, PT.
13. Wijngaards, N., Kempen, M., Smit, A. and Nieuwenhuis, K. (2006) Towards sustained team effectiveness, *Proceedings of the COINS 2006*, Hakodate, JP.
14. Yao, X., Turoff, M. and Hiltz, R. (2010) A Field Trial of a Collaborative Online Scenario Creation System for

Emergency Management, *Proceedings of the ISCRAM 2010,* Seattle, WA.