# Fusion and Policy Relaxation in Patient Healthcare


**Lundy Lewis**
Engineering Dept.
Altusys Corp.
Manchester, NH USA
lewis@altusystems.com

**N Parameswaran**
Computer Science Dept.
U. of New South Wales
Sydney, Australia
paramesh@cse.unsw.edu.

**Pradeep Ray**
Computer Science Dept.
U. of New South Wales
Sydney, Australia
p.ray@unsw.edu.au



***Abstract* – We consider a community of agents in the domain of patient healthcare where the community is populated by personal assistant agents residing on mobile phones assisting the human patients and sever agents managing the services offered to the patients. In this scenario, we look at patient care polices, and interpret them in the context of crisis situations. Existing models dealing with behaviours of such systems use policies to determine actions in predefined and static situations but do not address the problems that arise when policies have to be interpreted in changed situations. We investigate techniques for policy compliant actions in dynamic situations particularly when multiple actors are involved in a problem solving process in a crisis situation.**


## 1 Introduction

In a world inhabited by multiple actors where resources are shared in their problem solving processes, polices are necessary to provide openness, guidance and scalability. Building policy-aware systems has been increasingly drawing attention of researchers where agents are built to be aware of policies imposed upon them while accessing control over shared resources. While this approach produces policy compliant agents, the strategy for policy compliance in moments of crisis has remained an unexplored issue. When in situations such as fire, flood, tsunami, or earthquake, it is arguable whether polices that are followed in normal situations are to be strictly followed or followed at all. However, not complying with policies altogether will lead to unacceptable resource uses and problem solving behaviours.

We investigate the problem of policy compliance in crisis situations. Whenever an agent's actions do not comply with the policy specifications, we quantify the deviation of states and behaviours in order to enable the agents or users of the agents to judge whether the deviation is acceptable in the domain of applications.

## 2 The Application Domain

In view of the growing use of IT in healthcare management, particularly the use of mobile phones in developing countries, both patients and health service providers need automated techniques to plan and monitor large scale medical services. The use of light weight agents on the mobile devices of the patients, and heavy weight agents on the servers of the providers will form a natural solution to this problem. This solution has a greater significance when the agents have to work obeying all the policy requirements of the government and other organizational agencies. Our approach investigates techniques and strategies for challenges found in this application domain and in particular pays attention to the role of policies while offering services in emergency or crisis situations, e.g. caused by a disaster when policies are often likely to be violated for justifiable reasons.

We consider a community of agents in the domain of patient healthcare where the community is populated by personal assistant agents residing on mobile phones assisting the human patients, and sever agents managing the services offered to the patients. In this scenario, we look at patient care polices, and interpret them in the context of crisis situations. Existing models dealing with behaviours of such systems use policies to determine actions in predefined and static situations but do not address the problems that arise when policies have to be interpreted in changed situations. We investigate techniques for policy compliant actions in dynamic situations particularly when multiple actors are involved in a problem solving process in a crisis situation.

## 3 The General Approach

In cooperative work, complying with the requirements of policies of the domain is necessary. However, it is also the case that policies do not necessarily always take into account all situations in a domain, particularly when the domain is subjected to a crisis situation. In such cases, any action that an agent may perform may violate the policy constraints, and the action may have to be aborted. Action abortion may not always be a good option and it thus becomes necessary to investigate if polices have to be violated, how they can be violated, and what strategies are available to an agent caught up in the crisis situation. The core idea that we propose in this paper is that the agent must examine its available options and choose to deviate carefully from the behaviours permitted by the policy, so that the chosen behaviour only becomes minimally noncompliant.

Fundamental to the notion of deviation is the notion of distance measures between concepts. Distance between concepts can be computed using similarity measures which are often obtained by mapping the concepts using the notion of ontology mapping. Ontology mapping is performed using different techniques, including heuristic based, semantic based, and statistic based approaches. For example, [1] is a statistical approach that uses machine learning techniques. In [2], we proposed a logic based approach to this problem. Recently in [3], we showed how semantics expressed in English from Wordnet can be used for computing the semantic deviation of one concept from the other using sentence simplification methods. Investigating the limitations of policy compliance requirements in crisis situations is thus an important aspect of cooperative problem solving in dynamic domains.

## 4 Innovations

The following innovations are introduced in this paper.

- Techniques for incorporating policy awareness in agents – We investigate techniques for making agents policy aware, focussing in particular patient care in the healthcare domain.
- Strategies for crisis management in the context of policy driven behaviours – We investigate techniques for relaxing or tightening policies in the presence of crisis.
- Techniques for quantifying semantic distances between behaviours – Modelling behaviour as a sequence of states, we devise methods for quantifying the semantic "distance" between behaviours extending available ontology mapping techniques.

We consider a typical application scenario where there are several patients each with a communicating device (for example, a mobile phone) and a set of services offered by a medical service centre. We assume that an intelligent light weight agent $c_i$ resides on each patient's mobile device, and a set of heavy duty agents $a_i$ reside on the servers at the service location with larger computing capabilities possessing sufficient amount knowledge required for the services they plan to offer. Let $C = \{c_1,\ldots,c_n\}$ be the set of all patient (client) agents and $A = \{a_1,\ldots,a_m\}$ be the set of all service agents. We also assume that the agents in A are organized as a group, they actively service the requests coming from the agents $c_i$ whereas the client agents are individual agents, do not interact amongst themselves and thus there is no organization amongst the agents in C. The agents in C and A are guided by a set of policies which influence their behaviours. Our primary aim is to investigate the impact of a crisis situation on the problem solving behaviours of this community.

## 5 Related Work

Uses of policies for security and network management applications have long been known and yet policy awareness in agents seems to have received little attention. Some progress has been reported in web service based applications, where policy awareness has been shown to enhance the provision of openness, and distributed and scalable information access. Further, the declarative nature of policies is claimed to allow transparency of information access [4]. However, it has been recognized that policy awareness is only a part of compliance management and that further research is required. Policies also find applications in context aware services where context is used to make services personalized, adaptive, and not only reactive but also proactive. User context data including location, presence, time and identity information, are often extended to include any relevant user and environment information (also historical data, service subscriptions, biographical information, intentions and desires) that may improve communication and the service quality [5]. Crisis management has been an important topic of research and discussion, and problems from this domain have been looked at, solutions proposed and prototype systems implemented.

## 6 Architectural Considerations

We have examined typical policies in the healthcare domain and believe that the following architectural features provide the foundational components for agents.

The policies generally deal with, among other things, the real world, real time, and human attitudes. For example, consider the policy statement: If it's an emergency, you'll be taken to the nearest Accident and Emergency Department. In order to interpret this statement, agents need to understand terms such as emergency, nearest, etc. Emergency has the notion of time and agent attitudes built into it, and nearest is a special concept. Based on our studies of the healthcare domain [6], and our implementation experience [7, 8], we believe the following architectural features will be required in each agent.

- Belief structure (B)
- Attitude biased Goal Generator (GG)
- Task Pool
- Time Line Structure (TLS)
- TLS Manager
- Context structure (CS)

Briefly, an agent, from the belief structure, generates goals (G) using the goal generator GG, binds an attitude to each goal, and saves the goals as tasks in the Task Pool. An attitude is a behaviour selector that selects a particular biased behaviour of an agent amongst several possible behaviours in a given situation [7]. The agent then selects some of the tasks from the Task Pool, inserts them in the TLS as plans, and executes them in order to achieve the goal.

TLS is a complex data structure built over the time line using mental objects such as actions, plans, goals, events, state of execution, intervals, and future, present, and past segments. These objects may be abstract or detailed depending upon the situation.

The TLS manager decides where to insert a given task on the TLS and which tasks or subtasks on the TLS can be removed. The removed tasks may be either stored in the Task Pool for later insertion into TLS, or deleted forever.

In order to perform each action, execute each plan, achieve each goal, the agent uses a context structure CS, in which the context knowledge is maintained. For example, when a crisis has affected the world and if the agent believes that a new plan is required, then this fact is made available in the context structure for the action, and that action may not be executed.

## 7 Policy Awareness

In order to build policy awareness in the agents of the community, we need to systematically interpret the policies given in English and model them as agent behaviours. In order to do this, we will first need ontology of the domain, which we will assume is available:

Policy constructs

- Rights (to do certain action)
- Prohibitions (from doing an action)
- Obligations (towards performing an action)
- Dispensations (from doing an action)

Speech Acts

- Delegation (transfer right to another agent)
- Request (for action or right),
- Revoke (removal of right)
- Cancel (cancel request)

As an example, consider the following policy rule:

Rule 1: A patient has the right, upon request, to be given the name of his/her physician.

This can be simplified and understood as follows:

A patient has the right to put in a request, and when he puts in a request, the service centre is obliged to provide the name of the physician involved.

We then build this semantics into the agents as follows:

Let $c_i$ be the patient agent, i.e. the agent that works on behalf of the human patient. Consider the situation when $c_i$ puts in a request to a service agent $a_j$. Then, $a_j$ receives the request and generates a commitment to provide a reply at a later time. This commitment is implemented by inserting an action $v_k$ in the group time line structure (group TLS) where the action $v_k$ when performed initiates an appropriate service to send the details of the involved physician to the patient agent. In general, $v_k$ may be a primitive action, an abstract action, a plan, or a goal.

The challenges we identify are as follows:

- Identify those set of concepts derived from the policies which are likely to affected by crisis situations in interesting ways at various levels. These concepts can be basic concepts, abstract, or composite concepts, and include them in the policy ontology. For example, those concepts that relate to time criticalness and space criticalness.
- Based on the concepts, identify a set of basic and composite or abstract behaviours of the agents. Note that the behaviour may be a single agent private behaviour, social behaviour, or a multi-agent behaviour.
- Finally, map a give policy to a set of behaviours.

A crisis typically affects the world directly and an agent's internal states indirectly. For example, an accident in the world may affect the plans of an agent. In situations where the crisis is caused by an agent $a_j$, the internal states of another agent $c_i$ can be affected directly. Thus, the agent $a_j$ saying to the agent $c_i$, "You have cancer!" may cause a crisis situation for $c_i$. In our research, we only address the former type of crisis situations.

In general, a crisis can affect an agent behaviour in several ways, and to perform a detailed study of this, we need to analyse how a given crisis situation affects the internal states of the agents in the context of achieving a policy p. Our preliminary examination of the healthcare policies suggests that an agent, during its course of execution, can face problems such as the following ones:

- Pushed into an unexpected state: If an agent is executing an action in a world state s expecting to go to the next state s′, the world may instead go to totally different state s″. That is, the agent may "think" that it has been pushed into an unexpected state.
- Pushed into an exceptional state: If an agent is executing an action, the world may go to an exceptional state.

In the above cases, it may be possible that no solution path may exist from the current state to the goal state, and policy compliant behaviour may not be possible. Problems such as the agent being not able to execute its next planned action,

We have seen how a community can interpret the policy Rule 1 above in a normal situation. However, this interpretation may not be appropriate always in a crisis situation, such as when the patient was treated in a temporary camp set up immediately after a tsunami by several paramedics and healthcare professionals, and the patient's request may not be served or may even be ignored for ever. An agent can achieve this by dropping a few "less important" goals from its Task Pool.)

Consider another example:

Rule 2: Except for emergencies, the physician must obtain informed consent [from the patient] prior to the start of any treatment.

The emergency above in the rule refers to the emergency the patient may be in. However, the service agents interpreting this policy may be in a crisis situation. For example, there may a situation where the doctors may have to treat too many patients that they may choose not to worry about the time consuming protocol of obtaining consents (for example, obtaining consents may involve making phone calls to insurance companies, signing forms, etc). An agent may achieve this by generating goals using goal generation rules that may have weaker preconditions.

In order to cope with crisis situations such as the ones mentioned above, agents need well defined strategies to respond to such challenges. As a solution to this, we propose to develop techniques for semantically relaxing the policies.

# 8 Semantic Relaxation of Policies

Semantic relaxation of a policy rule r refers to transforming the rule to a new rule r′ using new concepts related to the concepts originally used in the rules r. Thus, if H is a concept used in r, then r′ may use H′, where in the domain ontology H′ is the concept semantically closest to the concept H.

For example, consider the following policy rule:

- Rule: If you're pregnant and have served 12 month waiting period, you can enrol online.

Suppose that pregnancy is accompanied with diabetes, then we can replace enrol by call, and twelve month by zero days. Thus, the modified rule will look as follows:

- Modified Rule: If you are pregnant then you can enrol on line or call the emergency number.

We obtain the relaxed rule by deleting a term in the conjunct of the antecedent of the rule, providing alternate actions on the right hand side of the rule, where the alternate action was extracted from the domain ontology. We have assumed that for this example, the action call the emergency number is semantically closest to the action you can enrol on line.

The intuition behind this relaxation is that the new behaviour b′ that the new policy induces would be semantically similar to the behaviour b that the original rule itself would have induced.

While this idea may appear very challenging, our experience in computing compliance measure employing techniques analogous to semantic relaxation in the domain of information security has given us encouraging results[8].

We propose to augment the above relaxation technique with the following lower level heuristics:

- <u>Continue</u>: In a crisis situation, if an agent finds itself having been "pushed" from its current state s to a new state s′, the agent can continue the transitions from the new state s′ to the goal state.
- <u>Recover from exception</u>: But, if the agent finds itself in an exceptional state e after a crisis has affected the agent, the agent can continue its transitions to another exceptional state e′ that is closest to a normal state s, and from s, the agent can continue its transitions to the goal state.
- <u>Postpone the goal</u>: The agent postpones the goal later sometime. This can be view as a primitive form of temporal relaxation. We may do this more rigorously using a temporal ontology.

- Drop the goal: The goal is dropped without being achieved.
- Achieve the goal somehow: An attempt is made to achieve the goal in a policy noncompliant way with minimal policy violations.
- Achieve related goal: A related goal will be achieved.
- Partial achievement of goal: Achieve only some subgoals.

## 9 Summary and Outlook

Policy compliance in crisis management is of high priority in all countries. Healthcare organizations are actively promoting information technology solutions to address this problem. The approach described herein will leverage the multidisciplinary collaboration in intelligent modelling, artificial intelligence, computer science, public health and management science to address the problem of policy formulation in safeguarding the humans from risks caused by manmade and natural disasters. The approach will shows promise for training practitioners on this new and emerging area of policy awareness management for crisis situations.

While this project is of great benefit to domestic situations, it is equally valuable in the context of s international initiatives on policy awareness, e.g. the international consortium on e-Health sponsored by the International Telecommunication Union (ITU) and the World Health Organization (WHO). The approach may lead to potent information technology solutions for global policy compliance.

In future work we plan to provide more mathematical foundation and examples. The heuristics discussed in Section 8 are important; we plan to examine further evaluation of such rules using real-life or simulated scenarios. The work should be considered as a work-in-progress, and further results are forthcoming.

## References


[1] A. Doan, P. Madhavan, and A. Halevy. *Learning to Map between Ontologies on the Semantic Web*. In Proceedings of the 11th International Conference on the World Wide Web, 2002.

[2] A. Wong, P. Ray, N. Parameswaran and J. Strassner, *Ontology mapping for the interoperability problem in network management*. IEEE Journal on Selected Areas in Communications, Vol. 23, No. 10, 2005.

[3] L. Lu and N Parameswaran, *Sentence Simplification Based Ontology Mapping*. FLAIRS 2008.

[4] V. Kolovski, Y. Katz, J. Hendler, D. Weitzner, and T. Berners-Lee, *Towards a Policy-Aware Web*, WWW2005 Workshop on Policy Management for the Web, 2005.

[5] S. Bessler and J. Zeiss. *Semantic modelling of policies for context aware services*. WWRF 2006,.

[6] F. Daneshgar, D. Bunker, and K. Lee. *Identifying opportunities for enhancing collaboration in the Australian health insurance sector*. International Journal of Business Process Integration and Management", vol. 1, 2006, pp. 210 - 218.

[7] S. Au, and N. Parameswaran. *Attitudes for agents in dynamic worlds*, FLAIRS 2001, Key West, Florida, 2001.

[8] F. Yip. *A Framework for Semantic Enabled Applications: A Case Study in Regulatory Compliance*. Ph.D Thesis, School of Computer Science and Engineering, University of New South Wales, Sydney, 2010.