

Using Objects and Patterns for Building Compliance Agents in Healthcare

Submitted to the OOPSLA-98 Midyear Workshop on Healthcare

James P. Davis

HealthMagic Corporate Technology Center
1901 Main Street, Suite 1570
Columbia, SC 29201 USA
jimd@healthmagic.com

Michael N. Huhns
Ronald D. Bonnell

Center for Information Technology
University of South Carolina
Columbia, SC 29208 USA
huhns@sc.edu
bonnell@sc.edu

Abstract

This position paper presents preliminary results on defining a method and architecture for building agents that monitor compliance to specified plans or protocols in the healthcare domain. In the paper, we discuss an architectural pattern for a class of autonomous agents known as *compliance agents*, and the methods used to analyze the problem domain to identify the agents and select their problem-solving architecture. Compliance agents work to insure compliance to a set of prescribed guidelines. Object-oriented analysis methods are used to determine the behaviors of the agents. The principal observation about identifying agent behaviors is that compliance agents track “artifacts” of compliance through a lifecycle, and that deviations from prescribed compliance indicate opportunities for compliance agents to take remedial action. We use a number of graphical notations as a basis for articulating and exploring the problem-solving architecture for this system. Finally, we demonstrate how this methodology and architecture have been applied in a compliance domain for occupational safety in the healthcare environment.

1. Introduction

Intelligent software agents are becoming recognized as a preferable technique over more traditional, monolithic knowledge systems for constructing personalized, collaborative, distributed, intelligent applications in many domains (Huhns and Singh 1997). In healthcare, this approach is of particular interest (Szolovitz et al. 1994; Hayes-Roth et al. 1996) because of the increasing demand for individualization and personalization at the point where the healthcare system makes contact with people—including patients as well as caregivers. Agent-oriented systems techniques are relevant and useful in healthcare problems where there is benefit to the human in having personalized “advocates” acting on their individual behalf in the domain.

One of the key problem-solving tasks in healthcare is monitoring compliance to specified protocols and plan regimens. Protocols are operational guidelines defined for areas of practice, including prescribed treatment regimens (as in the case of patient disease management) and workplace occupational practices (as in the case of safety guidelines for healthcare workers). Our interest is in defining an analysis method and architectural pattern that will allow us to look across several different compliance domains to elucidate the underlying problem-solving task structure, and in devising an appropriate system architecture that will allow us to capture the appropriate domain knowledge to operationalize a cadre of personalized agents into service in different compliance monitoring domains in healthcare.

There are a number of application domains that involve compliance monitoring, where the performance of some activities must be monitored to insure that the impact or outcome of performing these tasks complies with guidelines set by a standards body. In a manufacturing enterprise, guidelines exist for tolerances on machined

parts that must be checked. In healthcare enterprises, it is up to healthcare workers to insure that appropriate workplace practices minimize the risks of transmission of disease and other health risks. In environmental management domains, guidelines stipulate limits on the impact that industrial processes can have on natural habitats. In all these cases and more, there is a general relationship that exists among the entity responsible for setting guidelines, the entity responsible for implementing specific procedures for adhering to guidelines, and the entity responsible for carrying out daily activities such that procedures defined for guideline compliance are followed.

These relationships can be depicted as shown in the stylized Dataflow diagram of Figure 1. We have three entities that participate in the guidelines compliance process. For discussion, we present this process for the domain of occupational safety in the healthcare workplace, where a regulatory agency such as OSHA is responsible for setting policy. An individual medical practice (such as a hospital, laboratory or doctor's office) is responsible for interpreting and devising specific compliance plans that are appropriate for the specific workplace. Individual healthcare workers are responsible for insuring that their daily activities are carried out in accordance with the compliance plans drafted for their workplace.

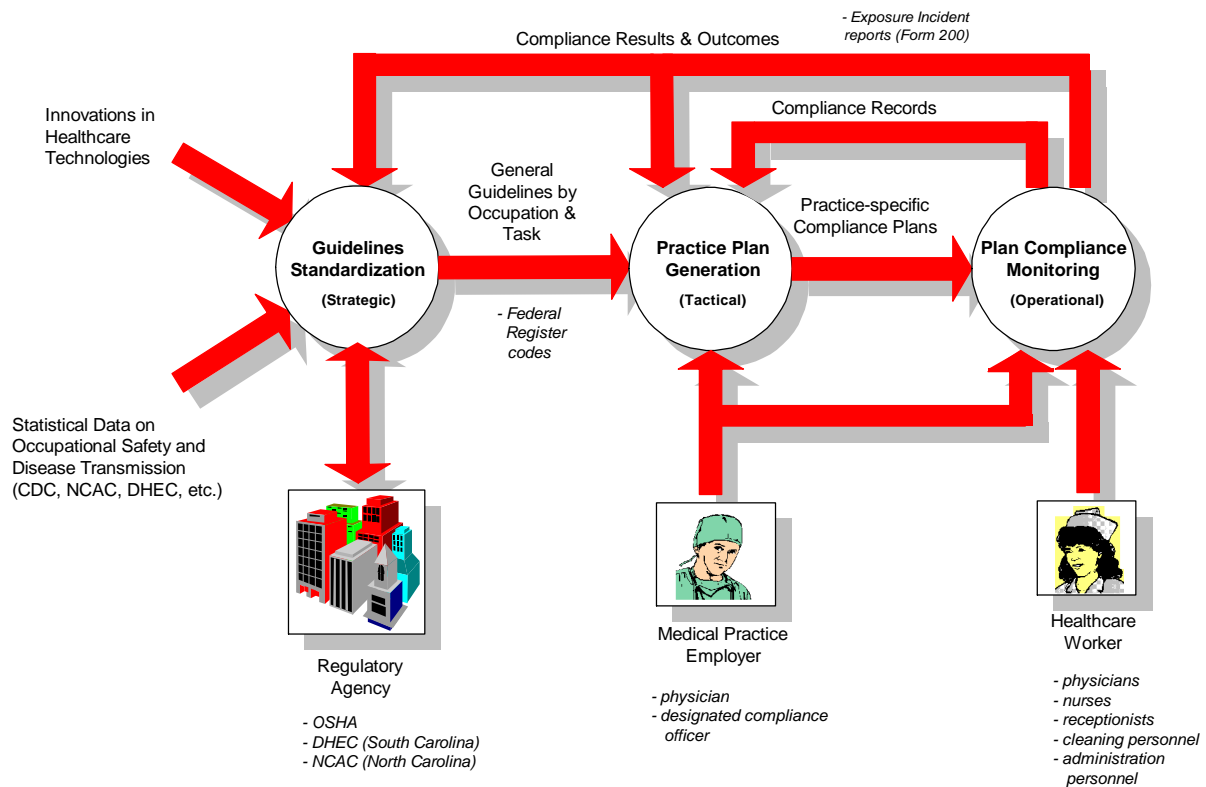


Figure 1. The compliance process for the occupational safety domain.

Today, the task of monitoring compliance falls on the shoulders of the healthcare workers themselves. In a typical medical practice, there is an employee who is designated as a “compliance officer” who has responsibility for drafting appropriate practice plans that meet the OSHA guidelines for workplace safety. The employees are responsible for following these compliance plans. The regulatory agencies, both federal and state, make periodic audits of medical practices to insure that the guidelines are being met.

The use of autonomous agents has been undertaken to enhance this environment for compliance. This added infrastructure addresses the following principal needs:

- Improve safety compliance through rigorous and continuous compliance monitoring;
- Support the workflow of protocol creation, refinement, and execution, thereby improving the collaboration among organizations that set guidelines, employers tasked with adhering to these guidelines through adoption of workplace-specific compliance plans, and the healthcare worker involved in daily activities related to compliance; and,

- Facilitate empirical feedback of aggregate data about compliance and risk in occupational populations to better manage the resources devoted to compliance monitoring.

In defining an architecture, we consider two principal problem areas. The first is how do we facilitate better management of the patient's treatment—in terms of compliance monitoring and ongoing assessment of a worker's progress against occupational guideline targets. The second is, if we can facilitate better occupational safety through automated monitoring of workers' progress against safety guidelines, how do we capture, store and manage these occupational protocols, allowing healthcare organizations to create, alter and "individualize" workplace-specific protocols for use with individual employees. The fact that these two problem areas are somewhat disconnected in their workflow allows us to partition the problem space into two primary subproblems: (1) protocol management, and (2) compliance monitoring.

At present, we have constructed an initial prototype for one component of the proposed architecture, the compliance system. Specifically, we have modeled in detail and have built a prototype for healthcare worker compliance monitoring for several components of a compliance plan for a medical practice, through the collaboration and interaction of software agents with both the medical practice and healthcare workers. The domain we have considered is that of compliance with OSHA occupational safety guidelines. The specific agent functionality we have explored and implemented is that of procedure compliance and associated follow-up for healthcare workers. As a result, we believe that software agent technology can be used as a means to deliver targeted and individualized functionality to carry out specific tasks associated with compliance monitoring and guidance for the healthcare worker, while also providing greater collaboration and better communication among the healthcare worker and medical-practice employer.

2. Methodology Overview

During the course of creating systems to monitor compliance using software agents, we have abstracted the essential process steps into a specific methodology for creating compliance agents. The process steps are defined as follows:

1. Partition and Decompose Problem Space. In the compliance domains we have examined, the process starts with the partitioning of the problem domain into strategic, tactical, and operational components. These flow in the manner shown in Figure 1. We have found that a high-level stylized Dataflow diagram (Edwards, 1993) is useful for visualizing this partitioning.

2. Identify and Create Inventory of Use Cases. A *use case* is a mechanism for identifying and describing the operational interactions between a system and the "actors" in its domain. We use the notation for use cases defined in the Unified Modeling Language (UML) standard (Rational, et al., 1997). We augment the notation by adding an icon to those use cases (ellipses) that likely contain one or more agents.

3. Identify and Localize Compliance Monitoring Activities. Given the inventory of use-case interactions, we refine each of these so as to (1) highlight the essential functional transformations of the interaction, (2) define the order and direction of interactions, and (3) identify the important "artifacts" of interaction, namely, the data storage elements required to process the interaction. We prefer to use the Dataflow notation for this purpose. At this time, we identify opportune locations in processes supporting the system interactions that might have autonomous agents associated with them.

4. Create an Inventory of Agents. Once the agents are identified relative to the system processes, we create an inventory, indicating the agent category name, the actions that trigger the agent to take action, the data sources the agent uses in its reasoning activities, the basic nature of its output actions, and the destination "actors" who receive some notification of compliance information. We use a tabular form to present this information.

5. Identify and Model the Essential Compliance Lifecycle. In our formulation of compliance problem solving, each compliance agent manages one or more "artifacts", each having a sequence of state transitions constituting a "Lifecycle". Our treatment of Lifecycle is based on that defined in the object-oriented analysis literature (Shlaer and Mellor, 1992). Each specific category of compliance agent has its own Lifecycle objects that it manages. We

have had good success in using the state transition diagram (STD) notation as an analysis and documentation medium for assessing the completeness of Lifecycle depiction.

6. Map the Abstract Architecture onto the Agent Environment. At this point, we need to map each compliance agent into the processing environment into which it will be delivered. In addition, we also need to indicate how the agents will interact with one another in that environment. It has been our experience that additional “helper” agents are needed in practice to facilitate distributed, coordinated agent behavior. To that end, we next model the sequences of interactions among the agents comprising the agent protocol. We use another notation component of the UML, sequence diagrams, for this model.

At this point, we have a reasonably complete model specification of the agents in their environment. This can be used as guidance for designing and implementing the actual agents in their environment. The method described herein can be used independent of implementation language or toolkit environment.

3. Architectural Overview

This section presents a high-level overview of the problem space, serving as the basis for creating a computing environment to manage compliance with safety guidelines through the use of protocol-driven autonomous software agents. We discuss basic assumptions about how such a system might be organized. We discuss the architecture by using diagrams of the principal components of the proposed infrastructure. These serve as an aid to succinctly capture the essential aspects of the system.

3.1. System Partitioning and Decomposition

There are two principal problem areas to be addressed in this architecture: (1) how to capture, manage and individualize workplace-specific treatment protocols; and (2) how to support workplace-specific compliance monitoring and treatment assessment to improve safety outcomes on an individual basis. We examine each of these problems separately.

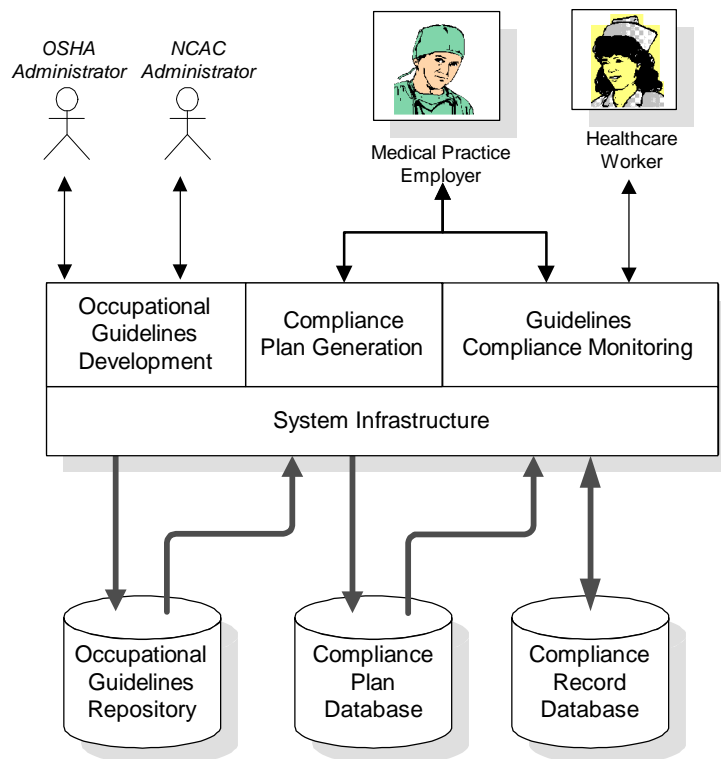


Figure 2. Partitioned system block diagram.

We organize the system architecture around three principal functions, as shown by the boxes in Figure 2. To further refine this description, we indicate the key sources of information that would either be generated or required to carry out the functions of each component. Finally, we indicate the principal users of each component.

Note that the Protocol Management enterprise depicted in Figure 1 is itself divided into two parts in Figure 2, namely Occupational Guidelines Development and Compliance Plan Generation. This reflects a natural delineation between features and functions to support the capture and management of the more general guidelines (which would likely be prescribed for use by both OSHA and state environmental and occupational regulatory agencies) and those to support defining compliance plans for a particular medical practice.

Overall, the system would function as follows, based on the diagram shown in Figure 2.

1. Guidelines Development - Users of this module would specify broad compliance guidelines for each area within the domain. These protocols are stored and indexed in a repository of protocols. They would likely be modified and extended over time, in light of new clinical results and discoveries.
2. Plan Generation - Users of this module (most likely, practicing physicians or compliance officers within a specific medical practice) would select guidelines from the database for purposes of constructing a specialized compliance plan for the healthcare workers in their workplace. This might entail selecting a cleaning regimen, specifying target values or ranges for measures reflecting levels of environmental contamination, or identifying specific treatments to be administered in the event of exposure to hazardous materials. Using the protocol as a guide, a plan is specifically drawn up for the healthcare workers and stored in the compliance plan database.
3. Compliance Monitoring - Users of this module (the practicing physician and healthcare workers in the medical practice) would be engaged by intelligent software agents in monitoring compliance to the generated guidelines. Software agents would execute on behalf of a healthcare worker. These agents would perform specialized tasks associated with (1) monitoring for activities that show compliance with components of the plan prepared for the medical practice; (2) checking for certain trends in occupational data collected and stored in the healthcare worker's summary compliance record that would place the healthcare worker outside of the guideline as represented by the compliance plan; and, (3) notifying, reminding and alerting either the healthcare worker, physician/compliance officer, or both when the agent(s) detect noncompliance.

The remainder of this paper focuses on laying out the detailed architecture and analysis of the compliance-monitoring functionality. The Guidelines Development and Plan Generation components are outside the scope of this paper, but have been presented to establish an appropriate context in which to discuss the structure and behavior of agents within this architecture. We believe that our architecture formulation is general enough to be used across different problem domains that fit the compliance problem statement presented here.

4. Detailed Architecture for Compliance Monitoring

Given the general architectural model we identified in the previous section, we now explore the detailed structure of the compliance-monitoring module.

4.1. Inventory of Agents in the Compliance Module

We have been investigating four categories of monitoring agents: (1) training compliance agent; (2) vaccination compliance agent; (3) exposure processing compliance agent; and, (4) personal equipment processing agent. Each was chosen to demonstrate a different type of processing using different types of data stored in the treatment plan and health record. There are also other types of agents that manage resources for the compliance agents. The training compliance and vaccination compliance agents monitor a patient's compliance to prescribed states in employment requirements. The agents compare the presence, or absence, of data in the compliance record against criteria as defined in the individual work plan for the healthcare worker.

The second two agents, however, address a different aspect of compliance management, the close monitoring of the worker's progress along specific process guidelines. This is based on bar code or other lab data used by the agents

to assess where the healthcare worker is in relation to their progression along a specific guideline. Table 1 provides a functional summary of this initial set of agents for this application domain.

Agent Category	Agent Trigger	Input Sources	Output Contents	Output Destination
1. Training Compliance Agent (TCA)	New healthcare employee is hired.	Training guidelines for new employees regarding handling of chemical and biohazard substances.	<i>Reminder</i> If new employee hasn't completed training within prescribed timeframe. <i>Alert</i> If maximum reminder count has been exceeded.	Reminder email sent to healthcare worker and echoed to physician. Alert email sent to employer.
2. Vaccination Compliance Agent (VCA)	New healthcare employee is hired.	Vaccination guidelines for new employees, as prescribed in the compliance plan.	<i>Reminder</i> If new worker hasn't had vaccination. <i>Alert</i> If maximum reminder count has been exceeded..	Reminder email to healthcare worker. Reminder notice echoed to physician. Alert email/page sent to physician.
3. Exposure Processing Agent (EPA)	Healthcare worker reported to have been exposed to a hazardous substance.	CDC treatment guidelines to be followed upon documented exposure incident, stored in compliance plan.	<i>Notification</i> or <i>Alert</i> to managing physician or compliance officer, based on nature of exposure and the type of guideline.	Notification page to managing physician or compliance officer. Ongoing progress via email.
4. Personal Equipment Agent (PEA)	Piece of clothing or personal equipment is issued to the worker.	PPE component of the compliance plan, which tracks issuance of items assigned to worker.	<i>Reminder</i> If worker hasn't completed equipment processing cycle within prescribed timeframe. <i>Alert</i> If maximum reminder count has been exceeded.	Notification or Alert paging message, depending on type of equipment and the processing guidelines dictated for the equipment.
5. Activity Monitoring Agent (AMA)	Triggered by the arrival of new compliance plan or compliance record entries.	Event signals provided by the system, generally the OS mechanism.	<i>Notification</i> to appropriate agents for individual healthcare worker.	Notification signal to compliance agents to "wake up" and evaluate any state changes.

Table 1. Summary of characteristics for each agent category.

4.2. General Transaction Process for Monitoring Agent Types

The figure below shows the principal transaction associated with each type of compliance monitoring agent. As indicated in Table 1, each category of monitoring agent uses its own unique information sources in the Compliance Record and Compliance Plan data structures for the healthcare worker. However, all categories share the same general pattern of behavior, described as follows.

The system provides some event triggering mechanism, indicating that new data has been placed into the Compliance Record or into the Compliance Plan for a specific healthcare worker. The type of record entry and content of the entry must be examined to determine whether the event is a type recognized and claimed by any of the agents currently attached to the worker. The activity monitoring agent (AMA) performs this checking activity. The AMA agent will send a message to "start" or "wake up" the appropriate compliance agent that might be interested in the new compliance record or compliance plan entry for its healthcare worker.

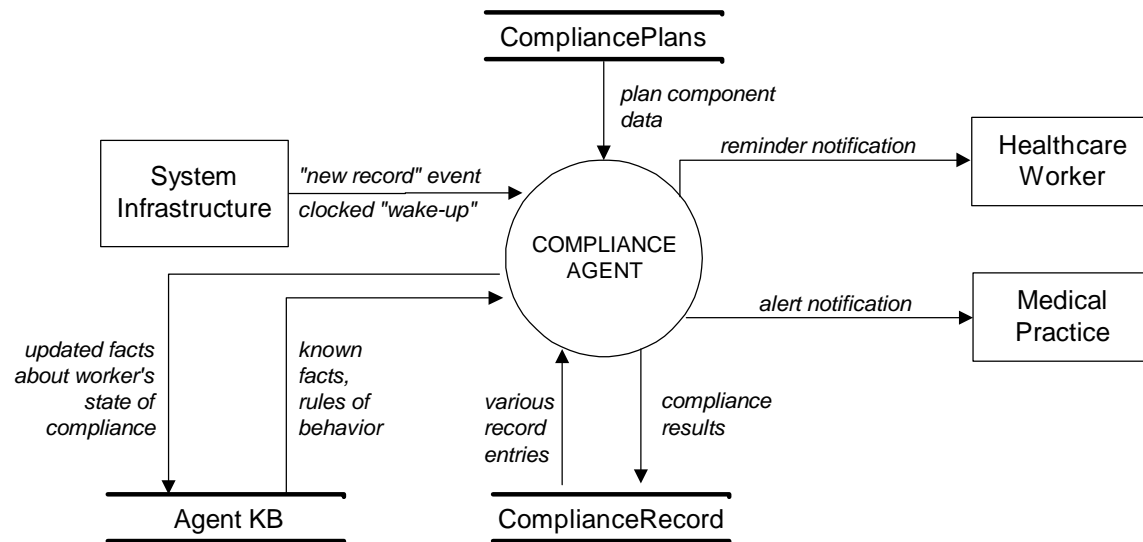


Figure 3. Essential data flow diagram for general form of compliance agent.

Once activated, the specific monitoring agent would evaluate the information for itself to see if it is of interest. The agent determines whether new data is of interest based on the type of compliance record or compliance plan entry. For example, a given worker's vaccination compliance agent (VCA) would be interested in entries in the Compliance Plan related to vaccination requirements for the workplace, created by the compliance officer through a system interface (which could be browser-based). It might also evaluate where the healthcare worker currently resides on the vaccination guideline, as defined by their specific Compliance Plan record. Once relevance is established, the agent would evaluate the data and compare it to target information contained in the compliance plan.

The agent retrieves its instructions, evaluates its possible actions, and selects its decision criteria from its own database of facts and rules. The facts component of an agent's knowledge base represents its current understanding of its sphere of interest and responsibility (generally, the state of the healthcare worker and his/her whereabouts on the particular component of the compliance plan). Its rules define the limits of what the agent can conclude from the data sources, and provide direction for its actions based on new information it derives as a result of "ruminating" over the facts it has stored.

As a result of the agent's actions, one of the following will happen: (1) new information is written into its internal facts base; (2) new information is written into the Compliance Record database; or (3) notification messages are initiated through an appropriate "mediating" interface (such as through a pager), or through corollary data generated for display on a dynamic page, whose URL is to be passed to a managing physician or compliance officer for the medical practice.

5. The Compliance Monitoring Agent Profile

This section of the paper discusses a typical compliance agent that has been studied and implemented in the current system. The section shows the remaining steps of our systematic methodology for analyzing and constructing an architecture for creating collaborating compliance agents. First, we present an overview of the agent's scope of functionality in the form of a use-case diagram. Next, we discuss the specific operational scenarios that the agent would handle.

5.1. Inventory of Transactions for Vaccination Compliance Agent

In the use case shown below, we indicate the individual transactions in which the VCA and other compliance agents participate, along with the other agents in the system with which the VCA interacts in performing these transactions.

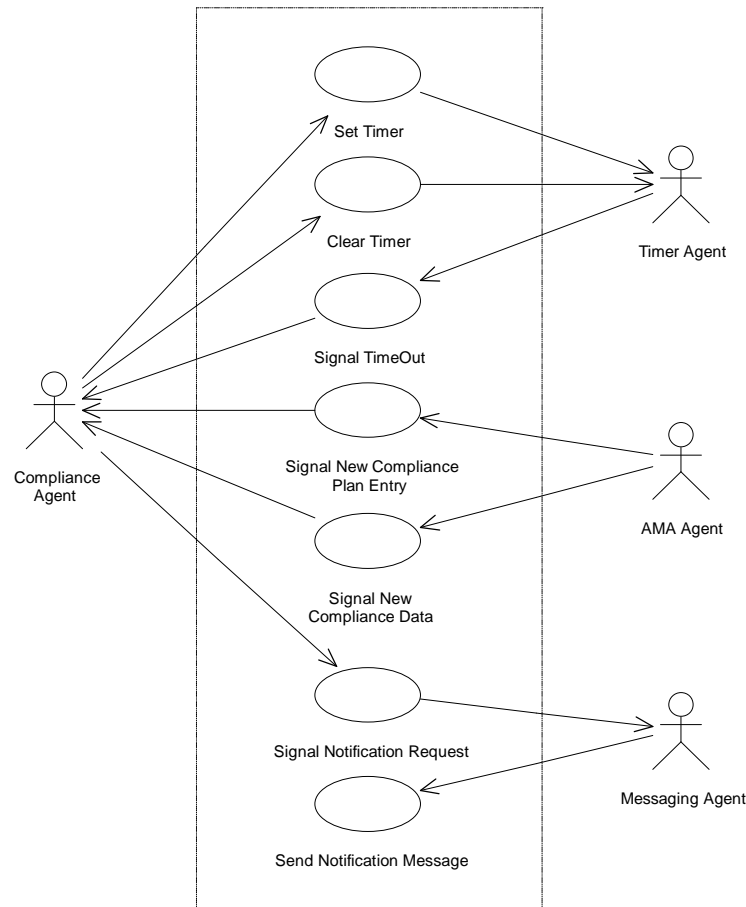


Figure 4. Use-case diagram for compliance agents.

The first three use-case transactions in the figure involve interactions between the VCA agent and a timer agent. The timer agent uses system timers to control the length of time that the VCA agent spends waiting for compliance input from the healthcare worker. This mechanism allows the VCA agent to take specific actions when anticipated data fails to arrive within prescribed time limits. This time limit is to be specified by the managing physician or compliance officer for the medical practice, and would be input into the healthcare worker's individualized compliance plan for each prescribed vaccination. Each of the following transactions—SetTimer, ClearTimer, and TimeOut—requires a “handshaking” protocol between the VCA and timer agents, the details of which are not presented in this paper. The next three use-case transactions involve interaction between the VCA agent and the AMA agent. The AMA agent is alerted to events occurring in the system environment by other application processes, and has responsibility of determining which agents have interest in these events, and forwarding appropriate information to these agents.

5.2. A State-Based Model for Prescribed Vaccination

During our analysis of the VCA agent, it was observed that the agent's actions were closely tied to the progression of a healthcare worker's given vaccination entry within the Compliance Plan. In other words, the agent was following a worker's compliance based on each vaccination that s/he is prescribed to take, and on where s/he is in the process of having it administered. In process modeling and object-oriented analysis, this is referred to as a “Lifecycle” model. The VCA agent manages a collection of PrescribedVaccination objects, one for each prescribed vaccination record written into the new healthcare worker's compliance plan.

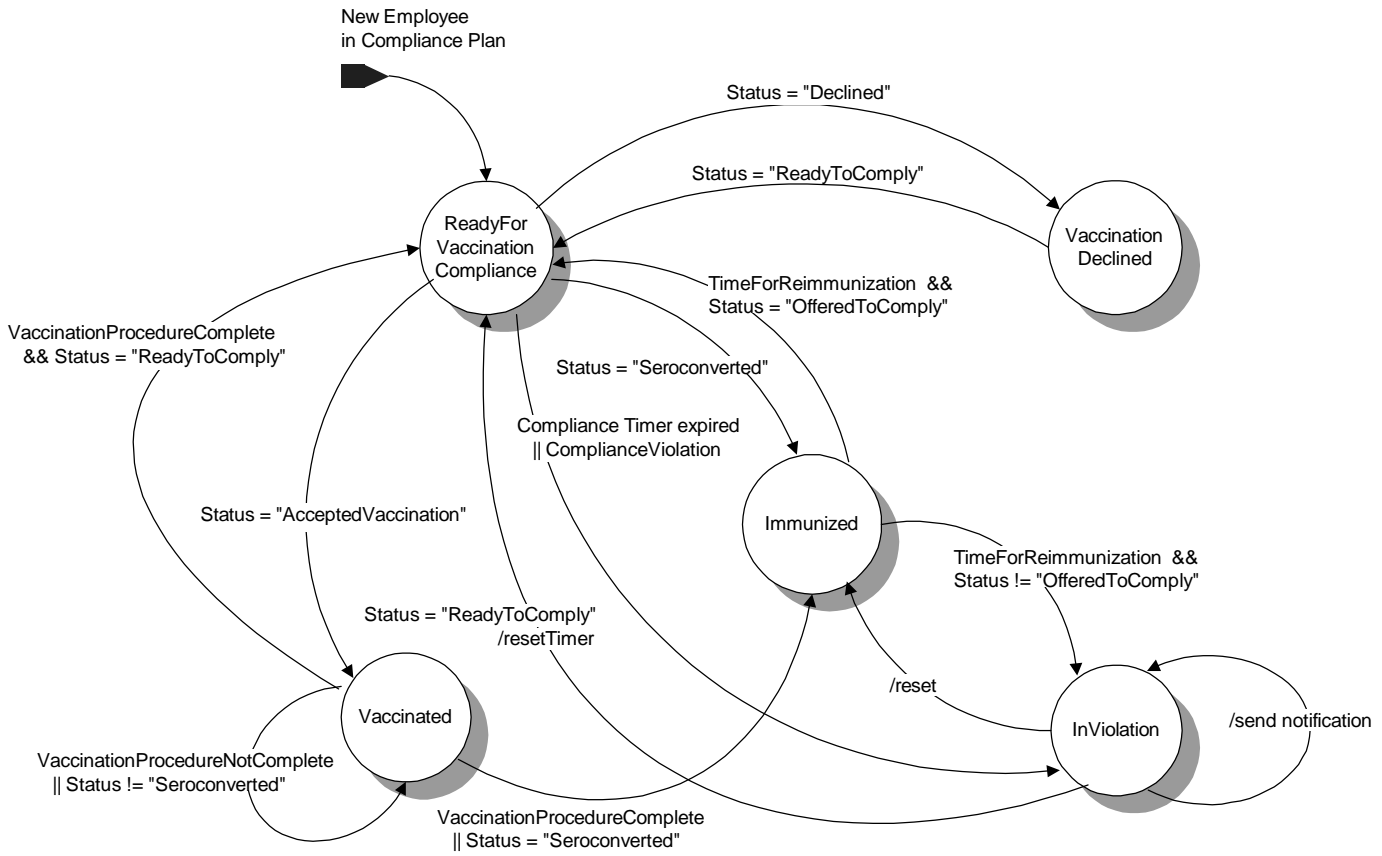


Figure 5. State model of vaccination compliance agent

The Lifecycle for the PrescribedVaccination object used by the VCA agent is shown in the figure above, depicted as a state transition diagram. When the managing physician or compliance officer for the medical practice prescribes a vaccination for each new employee and enters it as a record in a healthcare worker’s Compliance Plan, the event causes the AMA agent to signal the worker’s VCA agent. The VCA agent creates a new PrescribedVaccination object for this new Compliance Plan entry, placing the object in the initial “Ready” state. While in this state, the agent sets a timer to “remember” the vaccination time-out period specified in the Compliance Plan entry while it waits for a new entry in the healthcare worker’s Compliance Record indicating the vaccination has been performed.

5.3. Use Scenario for Vaccination Compliance Agent

The vaccination compliance agent performs several different activities to check the patient’s compliance to the prescribed regimen defined by the compliance guidelines for new healthcare workers. In this section, we examine one of these scenarios in detail and illustrate the use of a sequence diagram in the analysis.

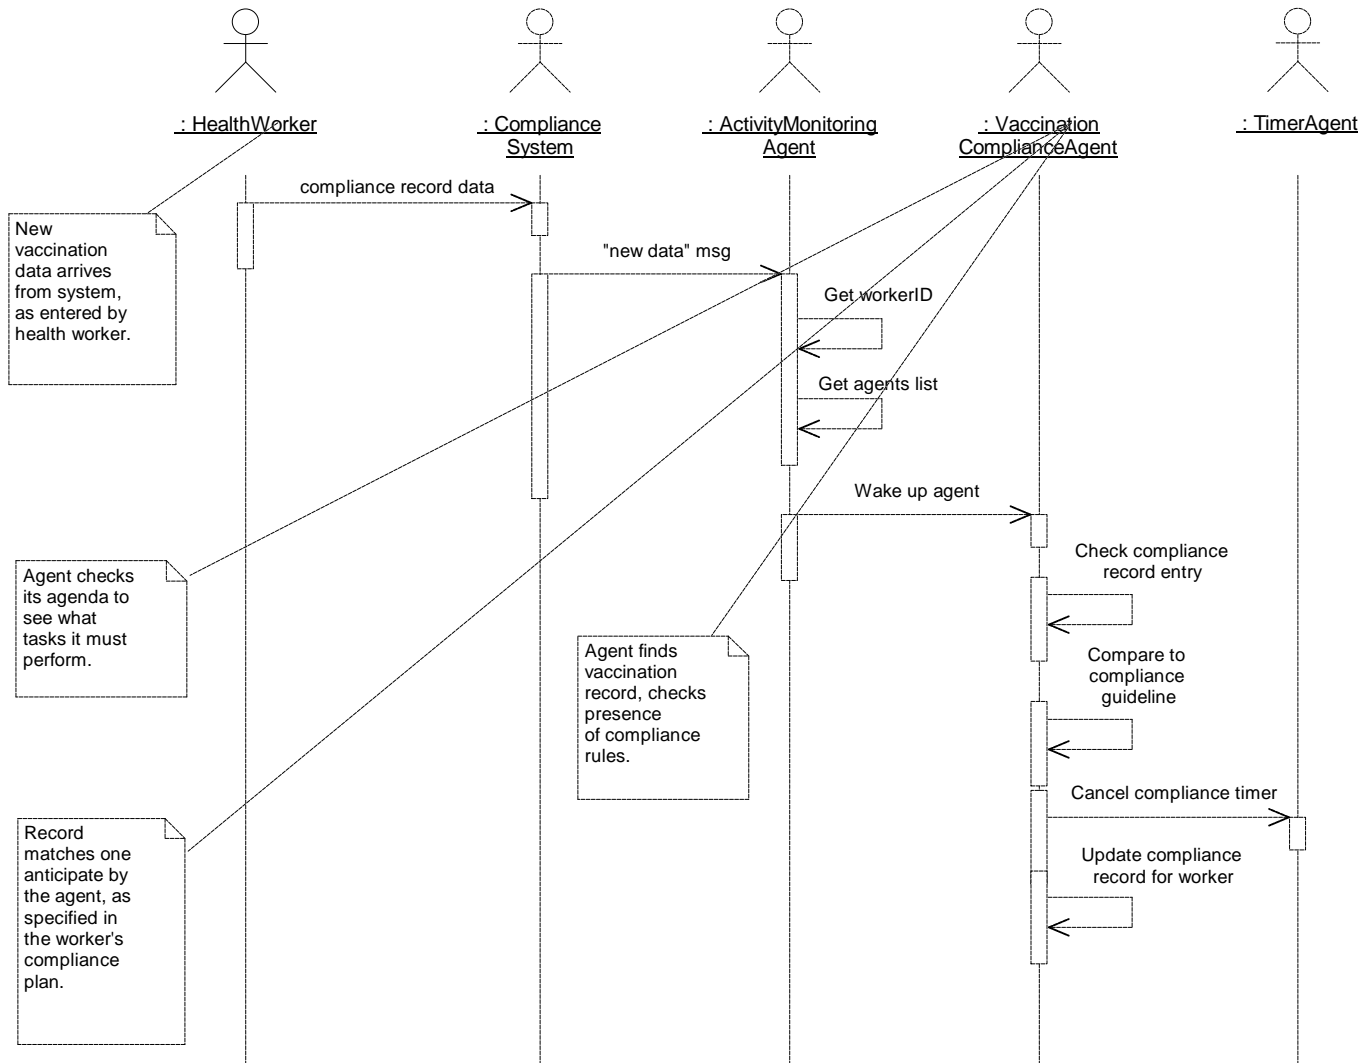


Figure 6. High-level sequence diagram for vaccination compliance checking.

A sequence diagram graphically depicts the ordered sequence of activities that transpire between various entities in the system. Along the top of these diagrams are listed the entities involved in the scenario. This usually includes one or more categories of intelligent agents. Time flows along the vertical lines dropping from the box representing each entity. The action on these vertical lines may involve a single entity or may involve one entity interacting with one or more other entities in the system. These interactions are represented by the horizontal lines with arrows and a text label indicating the nature of the communication.

5.3.1. Vaccination compliance checking

One of the key responsibilities of the vaccination compliance agent is to monitor vaccination compliance, in that the healthcare worker either has the appropriate vaccinations (such as for Hepatitis B). For compliance checking, as depicted using the Sequence diagram in Figure 6, a vaccination record is written into the worker's compliance record, which would also signal the activity monitoring agent that new record data has been written. The diagram shows that the data source provides a data record for the healthcare worker to the system, which writes the data and signals AMA. The AMA agent gets the worker's ID and determines the record type of new data, and uses this information to select the appropriate worker's agents to wake up. For new vaccination records, the AMA will select the patient's VCA to be activated. The VCA agent then attempts to match the new vaccination record with any of its currently active vaccinations for the patient. If it finds a match, it notes this in the healthcare worker's Compliance Record, then signals the AMA agent that it has completed its current task.

Once the healthcare worker's VCA is signaled by the AMA agent, it determines what tasks it needs to work on. It derives this task list from the state vector for the PrescribedVaccination object associated with the current prescribed vaccination. Information about the current state of the PrescribedVaccination object is stored in the Compliance Plan and in the Compliance Record for the worker. The Compliance Record maintains a persistent log of vaccination compliance actions that have been noted, based on prior tasks the agent has performed. In the current discussion, the VCA agent is interested in whether any new vaccination records for the patient correspond to specific ones it is anticipating.

The VCA agent uses the PrescribedVaccination object of each "active" prescribed vaccination entry in the Compliance Plan for matching against new vaccination records placed in the healthcare worker's record by the system. In this scenario, we assume that a new entry has been inserted into the record indicating a vaccination has been performed. Using rules in its knowledge base, the VCA agent finds a match and logs the event into the healthcare worker's Compliance Record. This completes the VCA's tasks for the current processing cycle, so it signals the AMA that it is finished. At this point, the MCA's computer process might either be put to sleep or destroyed (an implementation decision).

Finally, in the event of a match, the VCA agent writes an entry into the worker's Compliance Record, containing the following: (1) that a vaccination compliance event has occurred, (2) the identity of the vaccination (given either by its batch number and NDC drug code) and, (3) the compliance date.

5.3.2. Timer setting and time-out checking

At certain points, the AMA awakens the VCA to process one of the "time-out" events specified for the given vaccination in the Compliance Plan. A "time-out" is simply another type of event anticipated by the agent, set up around actions defined in the Compliance Plan, to allow the agent to take its own default actions when something it expects doesn't occur within a prescribed amount of time.

A separate timer agent carries out timer management. This agent handles initializing and resetting the timers. When a timer expires, the timer agent signals the AMA agent, which determines the appropriate agent to wake up. The AMA wakes up the VCA agent for a time-out event when no vaccination record has been written into the worker's Compliance Record within the allotted time frame that matches the anticipated vaccination from the Compliance Plan entry.

Given either scenario, the VCA agent responds by sending either a reminder or alert message indicating a compliance violation. For example, in the event of a vaccination time-out, a reminder message is first sent to the healthcare worker, instructing them to please complete the vaccination within a certain amount of time. On the other hand, if the "Reminder Count" threshold is exceeded, the VCA agent sends an alert message to the physician (or the assigned compliance officer in the medical practice) indicating that it has not received information on the vaccination. The physician or compliance officer may choose to take action by making a more direct inquiry in the healthcare worker's follow-up actions.

6. Summary

This position paper has presented our approach to using object-oriented methods and pattern-based architecture in the identification, definition and construction of software agent applications in the healthcare domain. We selected a narrow task-specific problem domain in which agents have high value, and analyzed the problem-solving structure for the compliance monitoring task, devising an appropriate architecture fitting the results of our object-oriented analysis. We have used a number of OOA/OOD notations--most notably, the Unified Modeling Language--as a basis for articulating and exploring the problem-solving architectural pattern for these systems. Using these notations made it easier to look across multiple agent applications and find the common architectural pattern.

We think of the architecture in terms of logical and physical components. The logical, knowledge-level components consist of the conceptual structures and behaviors for the agents, their knowledge sources, goals and agendas, and interaction protocols. These are mapped onto a physical architecture consisting of the JATLite agent toolkit coupled with the JESS inference engine and JDBC interfaces to a backend database.

7. References

- Edwards K. 1993. *Real-Time Structured Methods—System Analysis*. New York: John Wiley.
- Huhns, M.N. and Singh, M.P., "Agents and Multiagent Systems: Themes, Approaches, and Challenges," in *Readings in Agents*, M.N. Huhns and M.P. Singh, eds., Morgan Kaufmann Publishers, Inc., San Francisco, CA, 1997.
- Kautz, H., et al.. "An Experiment in the Design of Software Agents," in *Readings in Agents* (M. N. Huhns and M. P. Singh, Eds.), Morgan Kaufmann Publ., pp. 125-130, 1997.
- Kendall, E.A., Malkoun, M. T. and Jiang, C. "The Application of Object-Oriented Analysis to Agent-Based Systems," *JOOP*, pp. 56-65, February 1997.
- Larsson, J. E., and Hayes-Roth, B. 1998. Guardian: An Intelligent Autonomous Agent for Medical Monitoring and Diagnosis. *IEEE Intelligent Systems*, 12:58-64.
- Linster, M., and Musen, M. A., 1992. Use of KADS to Create a Conceptual Model of the ONCOCIN Task. *Knowledge Acquisition* 4, 55-87.
- Musen, M. A. 1989. Conceptual Models of Interactive Knowledge Acquisition Tools. *Knowledge Acquisition* 1, 73-88.
- Pont. M.J. and Moreale, E. *Towards a Practical Methodology for Agent-Oriented Software Engineering with C++ and Java*, Report #96-33, Leicester University Department of Engineering, December 1996.
- Rational 1997. *UML Notation Guide*, Version 1.1.
- Sheth, A., Worah, D., Kochut, K., Miller, J., Zheng, K., Palaniswami, D. and Das, S. *The METEOR Workflow Management System and Its Use in Prototyping Significant Healthcare Applications*, Large Scale Distributed Information Systems Lab, University of Georgia, http://lstdis.cs.uga.edu/publications/Workflow/82_tepr.html
- Shlaer, S., and Mellor, S. J. 1992. *Object Lifecycles-Modeling the World in States*. Englewood Cliffs, NJ: Yourdon Press.
- Szolovitz, P., Doyle, J., Long, W. J., Kohane, I., Pauker, S. G., 1994. Guardian Angel: Patient-Centered Health Information Systems. TR-604, MIT Laboratory for Computer Science.
- Tu, S. W., Shahar, Y., Dawes, J., Winkles, J., Puerta, A., and Musen, M. A., 1992. A Problem-solving Model for Episodic Skeletal-plan Refinement. *Knowledge Acquisition* 4, 197-216.
- Wooldridge, M. and Jennings, N. R. "Pitfalls of Agent-Oriented Development," *AGENTS-98*, 1998.