A Distributed Adverse Drug Reaction Detection System Using Intelligent Agents with a Fuzzy Recognition-Primed Decision Model

Yanqing Ji,[1,†] Hao Ying,[1,*] John Yen,[2,‡] Shizhuo Zhu,[2,§] Daniel C. Barth-Jones,[3,¶] Richard E. Miller,[4,‖] R. Michael Massanari[3#]
[1]*Department of Electrical and Computer Engineering, Wayne State University, Detroit, Michigan, USA*
[2]*School of Information Science and Technology, Pennsylvania State University, University Park, Pennsylvania, USA*
[3]*Department of Internal Medicine, Wayne State University, Detroit, Michigan, USA*
[4]*Veterans Affairs Medical Center, Detroit, Michigan, USA*

Discovering unknown adverse drug reactions (ADRs) in postmarketing surveillance as early as possible is highly desirable. Nevertheless, current postmarketing surveillance methods largely rely on spontaneous reports that suffer from serious underreporting, latency, and inconsistent reporting. Thus these methods are not ideal for rapidly identifying rare ADRs. The multiagent systems paradigm is an emerging and effective approach to tackling distributed problems, especially when data sources and knowledge are geographically located in different places and coordination and collaboration are necessary for decision making. In this article, we propose an active, multiagent framework for early detection of ADRs by utilizing electronic patient data distributed across many different sources and locations. In this framework, intelligent agents assist a team of experts based on the well-known human decision-making model called Recognition-Primed Decision (RPD). We generalize the RPD model to a fuzzy RPD model and utilize fuzzy logic technology to not only represent, interpret, and compute imprecise and subjective cues that are commonly encountered in the ADR problem but also to retrieve prior experiences by evaluating the extent of matching between the current situation and a past experience. We describe our preliminary multiagent system design and illustrate its potential benefits for assisting expert teams in early detection of previously unknown ADRs.

*Author to whom all correspondence should be addressed: e-mail: hao.ying@wayne.edu.
†e-mail: yanqing@wayne.edu.
‡e-mail: jyen@ist.psu.edu.
§e-mail: szhu@ist.psu.edu.
¶e-mail: dbjones@med.wayne.edu.
‖e-mail: remiller@med.wayne.edu.
#e-mail: mmassana@med.wayne.edu.

INTERNATIONAL JOURNAL OF INTELLIGENT SYSTEMS, VOL. 22, 827–845 (2007)
 • DOI 10.1002/int.20230

1. INTRODUCTION

1.1. Adverse Drug Reactions

Every day numerous patients suffer and even die of different types of adverse drug reactions (ADRs; i.e., drug-associated adverse incidents in which drugs are used at an appropriate dose). In 2000, for instance, there were about 100,000 deaths in the United States due to medical errors, of which about 7,000 were attributed to various types of ADRs.[1] These statistics do not include the number of ADRs that occur in nursing homes. It is additionally estimated that more than 350,000 ADRs occur in U.S. nursing homes each year.[2]

Even though premarketing clinical trials are required for all new drugs before they are approved for marketing, these trials are necessarily limited in size and duration and thus are not capable of detecting rare ADRs. Also, patients who participate in premarketing clinical trials often are typically a very limited subset of all the individuals who will receive the medication once it is licensed. Therefore, drug safety depends heavily on postmarketing surveillance—the surveillance of medicines once they have been marketed.

Current postmarketing surveillance methods largely rely on spontaneous reports, where health care professionals, drug manufacturers, and consumers are provided with Web-based forms (e.g., MedWatch™) by which they can report any suspected ADRs. However, there is a gross underreporting (perhaps as little as 1% of all cases).[3] Furthermore, the rate at which cases are reported is dependent on many factors, including the time since the drug launches, pharmacovigilance-related regulatory activity, media attention, and the indication for use of the drug.[4] In addition, as a passive system, latency and inconsistent reporting are limitations with less serious or unusual reactions.[5]

More recently, systematic methods for the detection of suspected safety problems (called signals) from spontaneous reports have been studied and practically implemented.[6,7] For example, calculation of the proportional reporting ratios (PRRs) has been incorporated into the routine surveillance activities of the U.K. Medicines Control Agency.[6] Another important signal detection strategy is known as the Bayesian Confidence Propagation Neural Network that has been used by Uppsala Monitoring Center in routine pharmacovigilance with its World Health Organization (WHO) database.[8] In the United States, the Food and Drug Administration (FDA) currently adopts a data mining algorithm called Multi-item Gamma Poissson Shrinker[7] for detecting potential signals from its spontaneous reports. Various methods such as empirical Bayes screening, reporting odds ratios, and incidence rate ratios have been used in other national (e.g., Australian) spontaneous reporting centers.[9,10] By utilizing data mining or Bayesian techniques, these methods can facilitate the evaluation of spontaneous reports to detect potential signals for further evaluation. However, the performance of these techniques could be highly situation dependent due to the weaknesses and potential biases inherent in spontaneous reporting. When confounding exists, the signals generated by a data mining algorithm may not represent a true causal effect relationship between ADRs and the drugs and thus are not reliable.

Because of their limitations, the current systems are not ideal for rapidly identifying rare ADRs. A more effective system is desirable. In this study, we propose an active multiagent framework that can potentially assist experts in the FDA for early detection of previously unknown ADRs by analyzing patients' electronic health data *in real-time* using available data sources distributed across different health organizations. In this framework, intelligent agents would actively assist safety evaluators at the regulatory agencies and physicians at hospitals and clinics, accelerating the process of postmarketing surveillance by assuring that the relevant information distributed across different locations can be gathered and utilized more expediently. Despite their distinctive features that make them especially attractive and advantageous for medical applications, to the best of our knowledge, there are no reports in the literature that employ intelligent agents or multiagent systems to address the ADR detection problem. Another novelty of this study is that, in this framework, intelligent agents assist a team of experts based on a well-known human decision-making model called Recognition-Primed Decision (RPD), a model proposed by Klein to describe how domain experts make decisions by recognizing the similarity between the current situation and past experiences.[11] We generalize the RPD model to a fuzzy-logic-based computational RPD model and integrate it into intelligent agents to help experts make complex decisions. In this fuzzy RPD model, fuzzy sets, fuzzy rules, fuzzy logic, and fuzzy reasoning are not only used to represent, interpret, and compute imprecise and subjective cues that are commonly encountered in this problem but also applied to retrieve prior experiences by evaluating the extent of matching between the current situation and a past experience.

The remainder of this article is organized as follows. We first review the existing research on multiagent technology and the RPD model. Then, a practical multiagent framework is proposed in Section 2. In Section 3, we discuss how a group of intelligent agents collaboratively detects previously unknown ADRs. We wrap up with conclusions in Section 4.

1.2. Background on Multiagent Technology

The primary feature of multiagent technology is the agents' ability to collaborate with each other. The agent communication language (ACL) enables the agents to communicate with each other at the knowledge level and unite their efforts to become a collective of working individuals, who are aware of each other's goals and intentions.[12] Therefore, ACL supports not only a set of standard performatives, but, more importantly, conversation protocols based on these performatives. These protocols define how an agent chooses its replies to a message. In an open environment such as the Internet, ACL must also support security, privacy, integrity of data, and authentication of an agent's identity. There are currently two widely accepted ACLs: KQML (Knowledge Query and Manipulation Language) and FIPA (The Foundation for Intelligent Physical Agents).

A multiagent system offers a natural way of tackling distributed problems, where each agent is a "smart" software program that acts on behalf of human users to find and filter information, negotiate for services, automate complex tasks, and

collaborate with other agents to solve complex problems. An important property of intelligent agents is their *autonomy*. Intelligent agents have a degree of control on their own actions, and under some circumstances, they are also able to make their own decisions, based on their knowledge and the information perceived from the outside environment. Thus, multiagent systems provide a good paradigm for those applications in which each component wants to keep its independence and autonomy from the rest the system. *Proactivity* is another important property of agents. Intelligent agents do not only passively react in response to external events but they also make goal-directed initiatives. For example, in a teamwork model (e.g., CAST[13]), agents in a team can anticipate information needs of teammates and proactively offer relevant information. *Sociability* is also a common characteristic of intelligent agents. In a multiagent system, an agent is able to interact with other agents to get desired information. Through complex communications, agents can negotiate and coordinate their actions and collaborate for accomplishing their task and achieving the goal of the system. The agents in a multiagent system are permitted to run on different computers that may be located in different places. Each agent may only have part of the information required to solve the problem.

The multiagent technology provides a new paradigm for developing software applications in a variety of fields such as industry, commerce, entertainment, and military.[14,15] In the health care domain, some exemplary applications include the AADCare (agent architecture for distributed medical care) system for integrating the patient management processes,[16] the Organ Transplant Management system for the management of organ and tissue transplants among different medical centers,[17] and the internal hospital tasks management systems for monitoring the application of medical protocols.[18] Taking advantage of the recent development of computer networking, intelligent agent paradigm, and decision-making methodologies as well as tools to handle deterministic uncertainty and subjectivity (e.g., fuzzy logic), multiagent systems seem be a promising approach for tackling the ADR problem, where data sources and knowledge are distributed and coordination and collaboration are required.[19,20]

1.3. Brief Introduction to the RPD Model

Decision models play important roles in helping human beings make complex decisions. Integrating a proper decision model into intelligent agents can make agents more active and encourage closer agent–human collaboration.[21] The classical decision-making approach,[22] conceptualized by Janis and Mann, thoroughly surveys all possible options, carefully weighs the costs, risks, and benefits of each one, ranks all of them based on current expected utilities, and then chooses the best one. This strategy is rational, quantitative, and suitable for novices. It can be unpractical if some information is missing in the process of decision making. In contrast to this classical optimizing-based decision approach, the RPD model[23] represents a naturalistic decision-making theory that employs a decision strategy called satisfying. Instead of trying to find the best solution, RPD identifies the first workable option based on previous decision experiences. This model is more qualitative, efficient, and suitable for experts. It represents human decision behavior better than optimizing methods and thus is widely accepted by decision makers.

Because our multiagent framework is designed to solve medical problems where experts' experiences are so important and the decision making is not a one-shot but an iterative process, the RPD model is a natural choice in our case. There have been several attempts at integrating RPD with agent technologies.[24,25] For instance, Norling et al. employed RPD in a BDI agent framework as a more realistic way of simulating human societies.[24] More recently, Yen et al. implemented an RPD-enabled team-oriented agent architecture (R-CAST) to support information sharing and collaboration among members of the team.[21,25] However, none of these implementations of computational RPD could handle subjective and imprecise information, and their action evaluation stage is either ignored or simplified. By computational RPD, we mean a quantitative and computable RPD model that is readily implantable by computer.

To overcome the above limitations, we have developed, in a systematic manner, a fuzzy logic-based *general-purpose* computational RPD model.[26] Fuzzy logic is a well-established methodology that is effective for systematic handling of *deterministic uncertainty* and *subjective* information.[27] Fuzzy set theory has been successfully used to solve challenging industrial and medical problems in practice,[28–31] some of which are very difficult to solve without it. Fuzzy logic provides an effective means for complex knowledge acquisition, interpretation, and representation. As a cognitive model, the RPD often involves subjective and imprecise information, and more importantly, the model itself exhibits uncertainty and ambiguity. Therefore, it is natural to use fuzzy set theory to formalize the RPD process, which makes it more practical, humanlike, and powerful.

2. PROPOSED MULTIAGENT SYSTEM ARCHITECTURE FOR ADVERSE DRUG REACTION DETECTION

The ADR detection problem could be more effectively solved by utilizing our proposed framework that takes advantage of the widely available electronic health data in different health organizations such as hospitals and insurance companies. Administrative coding of diagnoses and procedures is a routine part of billing procedures for medical care in the United States and is a universally available electronic data source that could be utilized in patient safety monitoring. These data are coded in the form of ICD-9-CM (*International Classification of Diseases*, Ninth Edition, *Clinical Modification*) and CPT codes (Physicians' Current Procedural Terminology), two coding standards widely used in the United States. The codes are very valuable because they provide direct or indirect evidence of the clinical state of the patient, comorbid conditions, and the progress of the patient during the hospitalization or visit.[32] Pharmacy data and clinical laboratory data are two other common electronic sources of coded data that may give indication of potential ADRs.

Below, we introduce the overall organization of our multiagent system and the functionalities of each intelligent agent as well as the fuzzy RPD model that will be incorporated into human assistant agents.

2.1. Multiagent Organization

There are six types of intelligent agents in this framework (Figure 1): the National Regulatory Authority Agent (NRAA), which has the highest authority and is responsible for necessary management and information collection, the Safety Evaluator Assistant Agent (SEAA), helping the safety evaluators in regulatory authority make decisions, Pharmacist Assistant Agent (MAA) and Epidemiologist Assistant Agent (EAA), utilizing the expertise of pharmacists and epidemiologists, respectively, and collaborating with SEAA in supporting the safety evaluator's decision making, Health Organization Agent (HOA), a broker and controller for each health organization, and the Physician Assistant Agent (PAA), providing an interface and helping a physician acquire useful information and make complex decisions. We anticipate that these intelligent agents would coordinate with each other to provide services to physicians at hospitals and safety evaluators at the regulatory agencies to accelerate the detection of unknown ADRs. However, because the organization of these agents will be a very important consideration in such a large distributed multiagent system, we propose a hybrid agent organization model (as shown in Figure 1) based on our careful analysis of the system's requirements and constraints (e.g., privacy, scalability, security, and generality).

Within each health organization, there would be a unique HOA and many PAAs (each assisting a unique physician). The HOA and all the PAAs take the

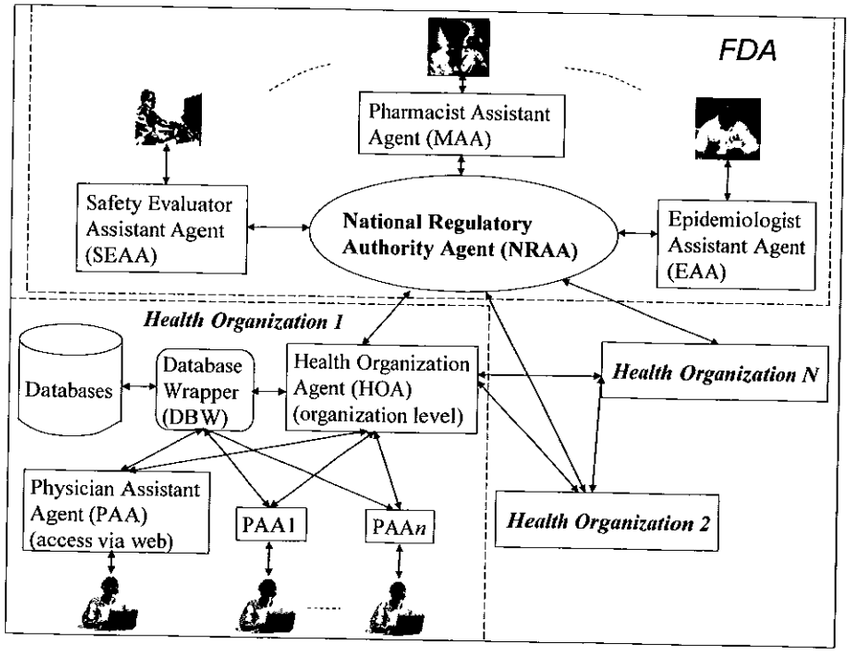

Figure 1. Hybrid agent organization.

classic Client/Server model, which is well known and widely used for distributed applications. In this model, although the server node provides most of the services or capabilities of the system, the client nodes just access and use them. In our case, the HOA works as a server that provides services for all the PAAs in the same health organization. The PAAs work as clients who can initiate requests. In the regulatory authority, the NRAA and all the expert assistant agents (e.g., SEAAs) also take the Client/Server model.

At the health organization level, a Client/Server model would not be desired because a HOA is expected to be an autonomous peer that can initiate communication and provide capabilities for other HOAs. One possible solution is to use the peer-to-peer (P2P) model, where all the peers are treated equally and the service is distributed among all the peers of the network. However, for the pure P2P model, it is difficult to maintain the coherence of the network and discover new peers. Also security is a quite important issue in this pure P2P model because each node is allowed to join the network without any control mechanism. We thus take the hybrid P2P model, a model between the Client/Server model and the pure P2P model. That is, a high level node (i.e., NRAA) is added to the top of all the HOAs, and at the same time different HOAs can still communicate with each other. To avoid being a possible communication bottleneck, the NRAA only provides necessary services (e.g., providing name–physical address associations for newcomers). Moreover, the caching mechanism can be implemented for each HOA so that the name–physical address associations are stored locally and the NRAA is contacted only when the address information is missing.

With this hybrid agent architecture, our system could offer two levels of detection for possible causality between a drug and an adverse effect. First, signal detection would be from a public health or epidemiological perspective by utilizing statistical or data mining schemes. This multiagent system has the potential to monitor all the available databases, and thus it could calculate population-based incidence rates (adjusted for multivariate confounding factors—e.g., age, gender, comorbidities, etc.) and provide epidemiological analysis (e.g., case-control and cohort analyses) to help safety evaluators make decisions. Second, signal detection could also be carried out at a single-case level by analyzing the details of key patient cases. Finding an unknown ADR is a complex process, and detailed analysis of the timing and plausible causality of ADRs often requires in-depth study of the sequence of events in individual patient cases. This system allows a safety evaluator to review the patients' administrative data gathered by his/her assistant agent, making the temporal relationship between the event and drug therapy more easily identified and organized.

2.2. Architectures of the Agents

National Regulatory Authority Agent (NRAA): This agent is at the top of the framework and has the authority to determine which HOA can join this framework and kills (malign) agents if necessary. It also collaborates with different HOAs in collecting useful information based on which safety evaluators in the regulatory

authority could make decisions. We anticipate that this agent would be managed by a third-party organization or national authority, for instance, the FDA. The NRAA has three important components:

- **Registration and Authentication Manager.** Each agent must be registered and authenticated through the Registration and Authentication Manager before it can be recognized by other agents in the framework. This component also provides a Naming Service that ensures that each agent in the platform has a unique name.
- **Team Manager.** Different experts in the FDA could form domain-based virtual teams to effectively communicate and collaborate with each other. This idea will be introduced in the next section. The functionality of the Team Manager is to manage the virtual team creation, evolvement, and deletion. It maintains a list of domains and the team members that joined in each domain. When a team member leaves its team, the Team Manager will enforce team membership so that future communications will not include the discontinued member.
- **Information Collector.** This component dynamically and continuously collects desired information based on the information needs of human experts in the regulatory authority.

Health Organization Agent (HOA): The databases in a health organization may not be directly accessed by outside agents due to the importance of privacy. Information can only be acquired through the corresponding HOA, which works as a broker and controls or filters the outgoing information (e.g., statistically deidentifying patient data, encryption). The architecture of the HOA and the PAA is shown in Figure 2, and the HOA has the following components:

- **Information Monitor.** This is one of the most important components in the HOA. It continuously monitors incoming and outgoing information through this HOA. It also

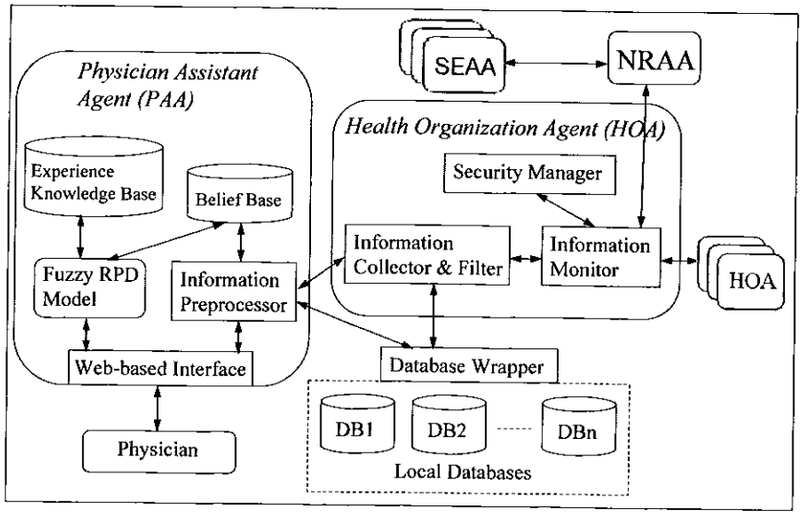

Figure 2. General architecture of a HOA and a PAA.

finds relevant new patient cases and collaborates in monitoring the further developments of the suspected patient cases.

- **Information Collector and Filter.** This component collects and filters information from databases, the PAA, and other HOAs. The raw patient information from databases has to be filtered and synthesized to get the useful information.
- **Security Manager.** This component takes care of security issues. Information transfer within the proposed system would comply with the Health Insurance Portability and Accountability Act (HIPAA) privacy regulations[33] and security standards, with only cryptographically secured communications occurring between agents, and (at least in the initial stages of investigation) only statistically deidentified information revealed to system users.

Physician Assistant Agent (PAA): This agent provides an interface through which a physician could more easily organize information about his/her patient and file online ADR reports. A PAA can directly access the electronic records of the physician's patient, but only through the corresponding HOA can it talk with outside agents (e.g., NRAA, other HOAs). We expect that each physician in health organizations has a unique PAA:

- **Information Preprocessor.** The Information Preprocessor processes the information gathered by the HOA before it is saved to the belief base or organized for the physician through a Web-based interface. For example, it could synthesize, organize, and summarize the returned clinical or pharmaceutical information about a patient in a more easily interpreted form (e.g., the calculation and graphical presentation of time intervals between critical events). It could also automatically extend the physician's description of a query so that a more comprehensive search could be performed.
- **Belief Base.** It represents the agent's view of the world. In the case of ADR detection, it maintains a set of cues that represent and characterize the current situation.
- **Experience Knowledge Base.** It keeps a set of experiences that were elicited from domain experts. Each experience contains four parts: relevant cues, expectancies, plausible goals, and courses of actions.
- **Fuzzy RPD Model.** The Fuzzy RPD Model could assess the current situation by comparing the cues (stored in the belief base) that the PAA perceives with the ones stored in the experience knowledge base and suggest possible actions to the PAA. For example, it could alert physicians regarding patients with potential side effects that may be worth reporting according to certain predefined experiences. This fuzzy RPD model will be introduced in the next section.

Assistant Agents at FDA: The detection of unknown ADR is a complex medical problem that needs expertise from different human experts (e.g., physicians, epidemiologists, and pharmacists). Each expert has a unique assistant agent to support his/her decision making. The SEAA preprocesses patient cases gathered from different health organizations and suggestions received from other human experts so as to provide easily accessible and useful information to the safety evaluators. Similar to the PAA, it also has a fuzzy RPD model that could assess the current situation and suggest plausible actions (e.g., give alerts on potential ADR-drug associations based on the agent's perceptions). The functionality of MAA and EAA is to assist their corresponding human experts to perform various analyses (e.g., epidemiologic studies) by gathering necessary information and providing a user-friendly interface.

Database Wrapper (DBW): This provides two important functions: (a) it provides a JDBC (Java Database Connectivity) wrapper library, a simple abstraction layer that encapsulates standard database language (e.g., SQL) and provides frequently used methods for database connectivity, working with database schema information, and working with database data, and (b) it offers a wrapper that can map from human description of the suspected patient cases to the different coding standards (ICD-9, CPT, etc.) because patients' electronic health data such as diagnosis are coded in the databases.

2.3. A Fuzzy Logic-Based Computational RPD Model

Instead of making abstract choices, the RPD model assumes that experts employ "situation-experience matching" decision rules in which they match a pattern of cues in the current situation with the corresponding cues in prior experiences. Once an experience is matched, the current situation is recognized as typical and four by-products are achieved: the types of goals, important cues, next expectancies, and a course of action that likely will succeed. In our proposed computational RPD model, fuzzy sets are employed to formalize the representation of imprecise cues, and fuzzy reasoning is used to abstract higher level cues from lower level elementary data. For example, to assess the causal relationship between a drug and an ADR, *temporal association* is an important cue whose value can be inferred from the time duration (t_d) between taking the drug and appearance of an adverse effect using the following rules:

If t_d is short, then *temporal association* is plausible
If t_d is medium, then *temporal association* is reasonable
If t_d is long, then *temporal association* is unlikely

where both t_d and *temporal association* are fuzzy variables and characterized by fuzzy sets.

A heterogeneous similarity measure is introduced to evaluate the degree of matching between the current situation and a prior experience. This similarity measure can handle different types of cues including *nominal*, *linear*, and *fuzzy*. We assume that the cues in current situation and the ones in a past experience have the same set of features or attributes, and the values given to them characterize the current situation and past experiences. Suppose that V and V' denote the set of cue values in the current situation and a prior experience, respectively, both of which are defined in feature space $U = \{C_1, C_2, \ldots, C_j, \ldots, C_n\}$. Then, V and V' can be represented by two vectors: $V = (c_1, c_2, \ldots, c_j, \ldots, c_n)$, $V' = (c'_1, c'_2, \ldots, c'_j, \ldots, c'_n)$ where c_j and c'_j could be *nominal*, *linear*, or *fuzzy* values. The similarity between V and V' is called *global similarity*, which is based on the *local similarity* between each cue-value pair c_j and c'_j. Below, we first propose a local distance measure and then map it to a local similarity measure. After getting a set of local similarities for each known cue-value pair, we then amalgamate them into a *global similarity* measure.

We define the following heterogeneous distance measure to calculate the local difference between c_j and c'_j:

$$d(c_j, c'_j) = \begin{cases} 1, & \text{if } c_j \text{ or } c'_j \text{ is unknown} \\ overlap(c_j, c'_j), & \text{if cue } j \text{ is nominal} \\ normalized\_diff(c_j, c'_j), & \text{if cue } j \text{ is linear} \\ fuzzy\_dis(c_j, c'_j), & \text{if cue } j \text{ is fuzzy value} \end{cases} \quad (1)$$

where the function of *overlap* is defined as

$$overlap(c_j, c'_j) = \begin{cases} 0, & \text{if } c_j = c'_j \\ 1, & \text{otherwise} \end{cases} \quad (2)$$

That is, for nominal cues, the cue distance is set to 0 if the cue values are equal; otherwise it is set to 1. The function of *normalized_diff* is defined as

$$normalized\_diff(c_j, c'_j) = \frac{|c_j - c'_j|}{\Delta_j} \quad (3)$$

The difference between the two cue values is used to represent the distance if the cue is linear. Δ_j is employed to normalize the cue values and is defined as

$$\Delta_j = a_j - b_j \quad (4)$$

where a_j and b_j are the maximum and minimum values for cue j, respectively.

Sometimes c_j and c'_j are not crisp or binary, and they should be evaluated by degree instead of hard boundary. If cues are defined in terms of fuzzy sets, then Hamming distance, one of the most commonly used distance functions, can be employed to define their distance. Let cue j be defined on the universe of discourse X and x be a generic element of X. If X is a collection of discrete objects, the function of *fuzzy_dis* is defined as

$$fuzzy\_dis(c_j, c'_j) = \frac{1}{m} \sum_{k=1}^{m} |\mu_{c_j}(x_k) - \mu_{c'_j}(x_k)|, \quad x_k \in X \quad (5)$$

where $\mu_{c_j}(x)$ and $\mu_{c'_j}(x)$ are the membership functions for the fuzzy values c_j and c'_j, respectively. If X is an interval $[\alpha, \beta]$, the function is defined as

$$fuzzy\_dis(c_j, c'_j) = \frac{1}{\beta - \alpha} \int_{\alpha}^{\beta} [\mu_{c_j}(x) - \mu_{c'_j}(x)]\, dx \quad (6)$$

Once we have the distance measure, we use the following mapping between the distance measure and the similarity measure:

$$s(c_j, c'_j) = 1 - d(c_j, c'_j) \quad (7)$$

Equation (7) only defines the similarity between a cue-value pair c_j and c'_j, and thus is called the *local similarity* measure.

The *local* similarities of each cue can be summed to give a total weighted *global* similarity between two sets of cue values V and V':

$$s(V,V') = \begin{cases} 0, & \text{if there exists } j \in (1,n), w_j = 1 \text{ and } s(c_j,c'_j) < \delta \\ \dfrac{\sum_{j=1}^{n} w_j s(c_j,c'_j)}{\sum_{j=1}^{n} w_j}, & \text{otherwise} \end{cases} \tag{8}$$

where $\delta \in [0,1]$ is a tolerance threshold set by human experts, and w_j is the weight for cue j, which represents the relative significance of that cue. Note that the weight of a cue is between 0 and 1, and 1 is only assigned to those required cues in a cue set. In Equation (8), once the *local* similarity of a required cue is less than a threshold, the *global* similarity will be set to 0 no matter how similar the other cues are.

When the degree of likeness between the current situation and a past experience is greater than a threshold according to the *global* similarity measure, the actions in the experience could be reused to solve the current problem. To select a workable course of action, we assume each action has a *condition* that is also characterized by a set of cues. Again, we use the *global* similarity measure to examine the matching between the values of the observed cues and the cues in a *condition* of an action. If all the matching succeeds in a course of action, then this course of action is a workable one.

We stress that the proposed *computational* RPD model is general and useful for many application domains, including medicine. Our attempt to implement a fuzzy-logic-based *computational* RPD model is a step toward practical RPD-enabled computer decision-support systems. The combination of this model with intelligent agents would greatly improve our capabilities to formalize and assist the decision-making process.

3. COLLABORATIVELY DETECTING UNKNOWN ADVERSE DRUG REACTIONS

In this section, we first outline how safety evaluators and their supporting human experts in regulatory authorities form different virtual teams according to the class of drugs they addressed. These collaborative agent teams could gather and synthesize diverse distributed data sources to present an integrated view to the safety evaluators. Then, we discuss the roles of different intelligent agents and how they collaboratively support the detection of previously unknown ADRs.

3.1. Virtual Team

Scalability is such a big issue in a large distributed system that it must be seriously considered. To reduce communication overhead, it is reasonable to divide people as well as their supporting agents into different groups or teams. For instance, a group of experts and agents may focus on a particular ADR. Due to the complex nature of ADR analysis, the team needs to allow new members to join (e.g., when a new interaction involving another drug is recognized) or leave.

We propose a dynamic virtual team strategy shown in Figure 3. A "domain" corresponds to a subarea in which different experts need to collaborate. For example, In the United States, well-trained safety evaluators in the FDA are assigned to review submitted ADR reports on a daily basis. Each safety evaluator is responsible for specific groups or classes of drugs. To further analyze the potential drug relatedness once the safety evaluator receives an ADR report, he/she may need support from other human experts such as pharmacists and epidemiologists. Thus, the safety evaluator and his/her supporting human experts can form a virtual team and collaborate in discovering a particular unknown ADR on a specific drug or drug class. Of course, a human expert and his supporting agent could be a member of multiple teams. For instance, there are two teams in Figure 3: *Team 1* consists of *Safety Evaluator A*, *Pharmacist A*, and *Epidemiologist A* as well as their supporting agents, which belong to *domain 1*. *Safety Evaluator B*, *Pharmacist B*, *Epidemiologist A*, and their assistant agents form *team 2* based on *domain 2*. Whereas *Epidemiologist A* is a member of both teams, *Epidemiologist B* does not belong to any team at present.

The formation of a virtual team would be transparent to the end users. When a safety evaluator makes a query or a human expert provides supporting

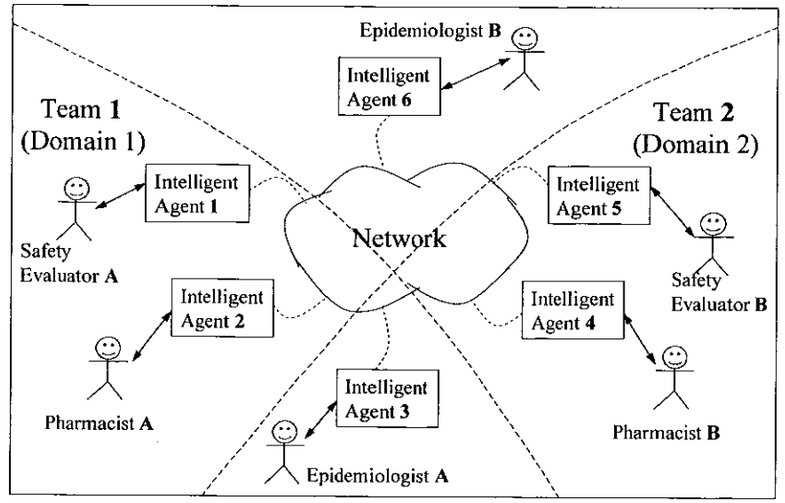

Figure 3. Domain-oriented virtual team.

information, he/she is automatically added to a team according to the suspicious drug (or its class) and the suspected adverse outcome. To achieve this, we must find out a strategy of classifying currently marketed drugs and embed the classification information into our proposed framework. The National Drug Code (NDC) System is chosen as a therapeutic or pharmacological classification scheme in which each NDC, a unique 10-digit number, serves as a universal product identifier for drugs reported to the FDA. In this coding system, the NDC Directory (NDCD) data were contained in seven ASCII Data Files, all of which have one or both linking elements: LISTING_SEQ_NO and FIRM_SEQ_NO. These two unique identification numbers are generated by the FDA and are only used in information systems to link tables together.

As shown in Figure 4, two files would be used to identify the class of a suspicious drug: *LISTINGS.TXT* and *DRUGCLAS.TXT*. When a NRAA receives a query request from a SEAA, it first extracts the drug name from this query. Next, the NRAA searches the *LISTINGS.TXT* file for the LISTING_SEQ_NO based on this drug name. Then, based on the returned LISTING_SEQ_NO, the NRAA queries the *DRUGCLAS.TXT* file and gets the PRODUCT_CLASS_NO. After that, the NRAA checks whether the team for this drug class already exists. If not, a new team is created and registered to the NRAA. Otherwise, the SEAA joins the existing team for this drug class.

To make the system more user friendly and to avoid information overloading the user, information delivered to the user will be divided into several levels for the user to prespecify. Also, when the human user makes a query through this system, he/she can choose which information level he/she is interested in retrieving/receiving. The system will also let the user choose the time span within

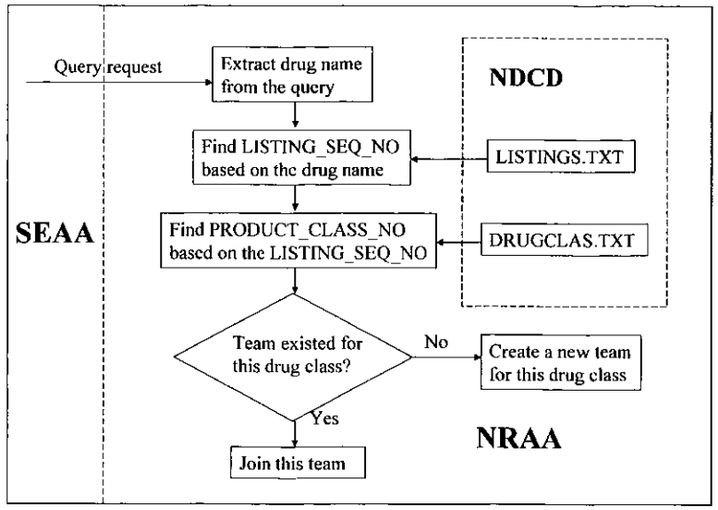

Figure 4. Virtual team formation and join.

which he/she is interested in the information and provide a GUI that supports rapid review and "drill-down" access to more detailed information.

3.2. Collaborative Detection Process

In the proposed agent system, each assistant agent adopts the fuzzy RPD model as a way to diagnose the current situation and assist its human expert in making decisions. Specifically, the fuzzy RPD model plays three important roles in this system. First, it helps agents find relevant experiences that could be reused to solve the current problem. Fuzzy logic is used in this process to determine the similarity or relevance between different pieces of information. Second, with its expectancy monitoring process, the RPD model can monitor the environment for unexpected events. Once these events are found, assistant agents alert relevant human experts, and at the same time, proactively provide information to the agents who potentially need it. Third, the fuzzy RPD model can help agents find out what information is missing in order to solve the current problem and track the further development of important information.

Physicians' timely reporting and insights are very important in postmarketing surveillance because they are the persons who are most familiar with their patients suffering from ADRs. A PAA could interactively prompt a physician for medical information specifically relevant to the particular suspected ADR to be reported. Thus more complete information could be collected, and filing a report would become simpler. A PAA could also help a physician take care of his patient by gathering information on previous patient encounters and pharmaceutical history for the patient. Furthermore, a PAA could continuously monitor a physician's patient cases and alert the physician about suspected side effects that might be worth reporting with the aid of the fuzzy RPD model (this functionality could improve the spontaneous reporting rate). The experiences used by the RPD model could be achieved by (systematic) solicitation with experienced experts or from relevant papers in the literature. For example, the WHO utilizes four cues (*temporal association*, *other explanations*, *dechallenge*, and *rechallenge*) to categorize the degree of causality association between a drug and an adverse event[34] based on an individual patient case. Among these four cues, *temporal association* is the most important one. It refers to the temporal relationship between taking the drug and occurrence of the adverse event. *Other explanations* denote alternative explanations by concurrent disease or other drugs. *Dechallenge* is defined as the relationship between withdrawal of the drug and abatement of the adverse effect. Contrary to *dechallenge*, *rechallenge* describes the relationship between reintroduction of the drug and recurrence of the adverse event. According to this classification scheme, a particular pattern of cue values characterizes a specific type of causality that may require certain courses of action to handle the ADR. Therefore, we can define various experiences, each of which is associated with a type of causality (e.g., certain, probable, possible). These experiences form an experience knowledge base that could be reused to assess the causality of suspected ADRs. The following is a sample experience:

Experience: The causality between a drug and a particular ADR is *certain*.
Cue: 1) The occurrence of ADR has a plausible temporal association with the drug.
2) There is *no* other explanations.
3) The ADR has *plausible* dechallenge relationship with the drug.
4) The ADR has *plausible* rechallenge relationship with the drug.
Goal: 1) Find the causality of suspected ADR.
2) File a spontaneous ADR report.
Action: 1) File an ADR report online.
2) File an ADR report through letter.
Expectancy: 1) Relevant information for filing an ADR report is available.
2) ADR report is successfully filed.

Note that the sample experience illustrated above is represented in natural language for the purpose of easy understanding. It could be expressed in a more computational way.

A HOA is the entrance point through which all the databases in a health organization could be accessed by outside institutes (e.g., regulatory authorities). The HOAs could collaborate with PAAs or SEAAs for tracking patients' information because a patient's electronic health data could be distributed across more than one health organization. They also help collect statistically deidentified epidemiological and biostatistical data for the regulatory authority.

The safety evaluators in the FDA play important roles in the process of discovering unknown ADRs. Under the help of their assistant agents (SEAAs), it would be easier for them to gather the details of a particular patient case and statistical data for epidemiologic study populations of interest. Figure 5 shows how a SEAA communicates with other health organizations to find a series of similar cases regarding a suspected drug–ADR pair. First, the SEAA asks the safety evaluator to define the study population and the relevant variables for controlling for confounding and effect modification (e.g., age, gender, polypharmacy, comorbidities, etc.). After that, the SEAA forwards the query request to the NRAA (path 2 in the figure). Then, the NRAA forwards that query request to all the HOAs (path 3) and, in the meanwhile, subscribes future similar cases of this drug to them. Next, all the HOAs query their local databases in order to find relevant patients (paths 4 and 5). Afterward, the patient data found by all the HOAs is deidentified and forwarded to the NRAA (path 6), which then forwards these results to the SEAA (path 7). Finally, the SEAA combines different cases and presents an integrated view to the safety evaluator.

With this system, safety evaluators could have at least three ways to discover unknown ADRs, which are complementary one another. First, they could continue to use current data mining algorithms based on collected ADR reports. Second, safety evaluators could perform patient-specific ADR reviews on individual cases by analyzing the details (e.g., temporal association, concurrent medications, and important laboratory data) of relevant patient cases gathered by their assistant agents. Third, the SEAAs could also help safety evaluators perform epidemiologic analyses at population level by collecting similar cases distributed across all

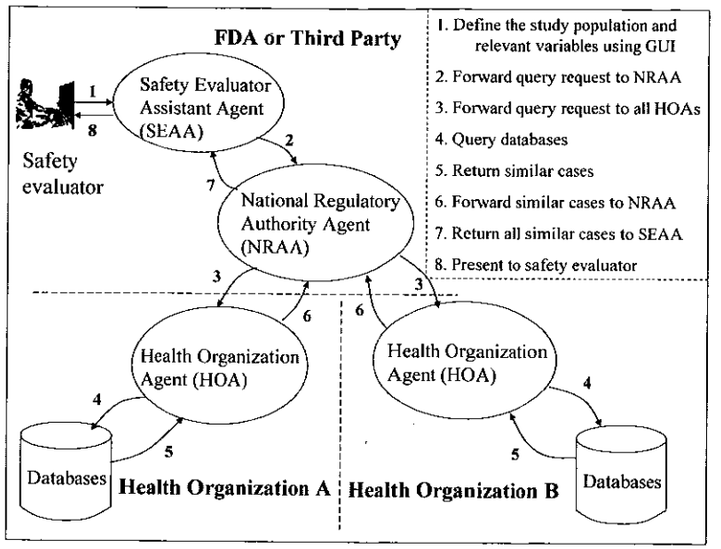

Figure 5. Collection of similar cases.

the health organizations and, in association with epidemiologists, assisting them to design epidemiologic studies (e.g., case-control or cohort studies). The SEAAs could also track the further development of the suspected patient cases. For example, a safety evaluator may be interested in what happens if the suspected drug is discontinued. With more information, safety evaluators could more easily determine the causal relationship between a new ADR and a specific drug.

4. CONCLUSION

The proposed multiagent framework is for detection of unknown ADRs that exhibits complexity, information distribution, and requirement of coordination. In this framework, a group of collaborative agents could effectively search and track relevant patient information, provide alerts on significant or unexpected events, and interact with a team of drug safety experts in the FDA on refining various suspicions related to an ADR of interest. Furthermore, the incorporation of a fuzzy RPD model makes an intelligent agent more effectively handle subjective and imprecise information and take over those decision-making tasks where human expertise is well developed. In summary, our proposed system provides a novel and systematic approach to helping regulatory authorities (e.g., the FDA) monitor the safety of marketed drugs by early identification of new and serious ADRs. Some important issues such as security and privacy have not been fully addressed in this article due to space limitations. They will be dealt with in our future reports.

Acknowledgment

This work was supported in part by a Wayne State University Research Enhancement Program grant.

References

1. Kohn LT, Corrigan JM, Donaldson MS. To error is human: building a safer health system. Washington, DC: National Academy Press; 1999.
2. Gurwitz JH, Field TS, Avorn J, McCormick D, Jain S, Eckler M, Benser M, Edmondson AC, Bates DW. Incidence and preventability of adverse drug events in nursing homes. Am J Med 2000;109(2):87–94.
3. Demiris G, Patrick TB, Austin Boren S. Assessing patient safety awareness and needs in rural hospitals in one US state. Inform Prim Care 2004;12(3):157–162.
4. Hartmann K, Doser AK, Kuhn M. Postmarketing safety information: How useful are spontaneous reports? Pharmacoepidemiol Drug Saf 1999;8(Suppl 1):S65–71.
5. Biriell C, Edwards R. Reasons for reporting adverse drug reactions—Some thoughts based on an international review. Pharmacol Drug Saf 1997;(6):21–26.
6. Evans SJ, Waller PC, Davis S. Use of proportional reporting ratios (PRRs) for signal generation from spontaneous adverse drug reaction reports. Pharmacoepidemiol Drug Saf 2001;10(6):483–486.
7. Szarfman A, Tonning JM, Doraiswamy PM. Pharmacovigilance in the 21st century: New systematic tools for an old problem. Pharmacotherapy 2004;24(9):1099–1104.
8. Lindquist M, Edwards IR, Bate A, Fucik H, Nunes AM, Stahl M. From association to alert—A revised approach to international signal analysis. Pharmacoepidemiol Drug Saf 1999;1(8):15–25.
9. Purcell P, Barty S. Statistical techniques for signal generation: The Australian experience. Drug Saf 2002;25(6):415–421.
10. Heeley E, Wilton LV, Shakir SA. Automated signal generation in prescription-event monitoring. Drug Saf 2002;25(6):423–432.
11. Klein GA. Sources of power. Cambridge, MA: The MIT Press; 1998.
12. Singh MP. Agent communication languages: Rethinking the principles. IEEE Comput 1998;31(12):40–47.
13. Yen J, Yin J, Ioerger TR, Miller MS, Xu D, Volz RA. CAST: Collaborative agents for simulating teamwork. In: Proc Seventh Int Joint Conf on Artificial Intelligence (IJCAI-01); 2001. pp 1135–1142.
14. Yen J, Fan X, Sun S, Hanratty T, Dumer J. Agents with shared mental models for enhancing team decision-makings. Decis Supp Syst 2006;41(3):634–653.
15. Parunak HVD. Applications of distributed artificial intelligence in industry. In: O'Hare G, Jennings N, editors. Foundations of Distributed Artificial Intelligence. New York: Wiley; 1996. pp 139–164.
16. Huang J, Jennings NR, Fox J. An agent-based approach to health care management. Int J Appl Artif Intell 1995;9(4):401–420.
17. Vazquez-Salceda J, Padget JA, Cortes U, Lopez-Navidad A, Caballero F. Formalizing an electronic institution for the distribution of human tissues. Artif Intell Med 2003;27(3):233–258.
18. Alsinet T, Ansotegui C, Bejar R, Fernandez C, Manya F. Automated monitoring of medical protocols: A secure and distributed architecture. Artif Intell Med 2003;27(3):367–392.
19. Ji Y, Ying H, Yen J, Fan X, Massanari RM, Barth-Jones DC. Team-based multi-agent system for early detection of adverse drug reactions in postmarketing surveillance. In: Proc 24th North American Fuzzy Information Processing Society, Ann Arbor, MI, June 22–25, 2005. pp 644–649.
20. Ji Y, Ying H, Barth-Jones DC, Yen J, Zhu S, Miller RE, Massanari RM. A team agent approach to postmarketing surveillance of adverse drug reactions. In: Proc 27th Annual Int Conf of IEEE Engineering in Medicine and Biology Society Shanhai, China, September 1–4, 2005. pp 6969–6972.
21. Fan X, Sun S, Mcneese M, Yen J. Extending recognition-primed decision model for human–agent collaboration. In: Proc Fourth Int Joint Conf on Autonomous Agents and Multi Agent Systems (AAMAS 2005); 2005. pp 945–952.
22. Janis IL, Mann L. Decision making: A psychological analysis of conflict, choice, and commitment. New York: Free Press; 1977.
23. Klein GA. A recognition-primed decision making model of rapid decision making. In: Klein GA, Orasanu J, Calderwood R, editors. Decision making in action: Models and methods. Westwood, CT: Greenwood Publishing Group; 1993. pp 138–147.
24. Norling E, Sonenberg L, Renquist R. Enhancing multi-agent based simulation with human-like decision making strategies. In: Moss S, Davidsson P, editors. Proc Second Int Workshop on Multi-Agent Based Simulation. Berlin: Springer Verlag; 2000. pp 214–228.
25. Yen J, Fan X, Sun S, McNeese M, Hall D. Supporting anti-terrorist analyst teams using agents with shared RPD process. In: Proc IEEE Int Conf on Computational Intelligence for Homeland Security and Personal Safety Venice, Italy, July 2004. pp 53–60.
26. Ji Y, Ying H, Massanari RM, Yen J. A fuzzy logic-based computational recognition-primed decision model. In: Proc 11th Information Processing and Management of Uncertainty (IPMU) in Knowledge-Based Systems Int Conf, Paris, France, July 2–7, 2006. pp 1660–1667.
27. Zadeh LA. Fuzzy sets. Inform Control 1965;8:338–353.
28. Barro S, Marin R. Fuzzy logic in medicine. Heidelberg: Physica-Verlag; 2002.
29. Ying H. Fuzzy control and modeling: Analytical foundations and applications. Piscataway, NJ: IEEE Press; 2000.
30. Mordeson JN, Malik DS, Cheng S-C, Malik DS, Cheng S. Fuzzy mathematics in medicine. Berlin: Springer-Verlag Telos; 2000.
31. Szczepaniak PS, Lisboa PJG, Kacprzyk J. Fuzzy systems in medicine. Heidelberg: Physica-Verlag; 2000.
32. Bates DW, Evans RS, Murff H, Stetson PD, Pizziferri L, Hripcsak G. Detecting adverse events using information technology. J Am Med Inform Assoc 2003;10(2):115–128.
33. U.S. Department of Health and Human Services. Standards for Privacy of Individually Identifiable Health Information. 45 CFR Parts 160 and 164; 2002.
34. Edwards IR, Aronson JK. Adverse drug reactions: Definitions, diagnosis, and management. Lancet 2000;356(9237):1255–1259.