in: Proceedings International Conference on Building and Sharing
of Very-Large Scale Knowledge Bases '93 (KBKS '93),
Tokyo, Japan, December 1–4, 1993, pp. 157–166

# Towards the Knowledgeable Community

**Toyoaki Nishida and Hideaki Takeda**
Graduate School of Information Science
Nara Institute of Science and Technology (NAIST)
8916-5, Takayama, Ikoma, Nara 630-01, Japan
nishida@is.aist-nara.ac.jp and takeda@is.aist-nara.ac.jp

## Abstract

The Knowledgeable Community is a framework of knowledge sharing and reuse based on a multi-agent architecture. In this paper, we describe the scope and goal of the Knowledgeable Community project and discuss how sharing of knowledge is achieved in the Knowledgeable Community. In particular, we focus on the organizational structure that facilitates mediation between those agents requesting for a service and those providing the service. We describe several techniques of mediation. Preliminary implementation of the Knowledgeable Community is in progress.

## 1 Introduction

Recently, importance of infrastructure has become one of critical issues of AI. Previously, two rather separate approaches have been taken to this end. One is to take a rather straightforward approach to building a large scale knowledge base system. The Cyc project [?, ?] is the most prominent one in this direction. Emphasis is placed on comprehensive collection of commonsense knowledge in a formal knowledge representation language based on first order logic. In contrast, the other approach is to develop a framework of knowledge sharing and reuse. Development of common languages and ontologies among interacting agents are major issues in the latter approach. The DARPA knowledge sharing initiative [?, ?] is the most representative approach in this direction.

These two approaches are considered to be addressing complementary aspects. It is worth seeking a way of integrating the two.

One way of integration could be regarding a large scale knowledge base as one of shared knowledge bases. However, this is not a good way to go. First, this would result in undesirable duplication of ontological definitions. Second, it is hard to understand and maintain a single very complex knowledge base.

In this paper, we take an alternative approach and explore a way of distributing large scale knowledge in a society of autonomous agents. The resulting framework, called the Knowledgeable Community, is characterized by ontology oriented agent organization and multi-stage mediation mechanism.

Ontology oriented agent organization means that agents are associated with *ontologies*, structural definitions of concepts. A mediation mechanism is provided for mediating agents looking for a service with those providing the service by combining concept similarity and definition. In order to derive even in the course of development, we propose building up the Knowledgeable Community in a bottom up manner.

Currently, implementation of the first version of the Knowledgeable Community is in progress using platforms developed by the DARPA knowledge sharing and reuse initiative [?, ?].

The rest of this paper is organized as follows. First, we discuss the importance of what we call *information-centered AI* as opposed to *processing-centered AI*. Second, we introduce the framework of the Knowledgeable Community, and feature design principles behind it. Third, we describe technical issues related to ontology-oriented agent organization and multi-stage mediation mechanism. Fourth, we discuss methodological issues on developing the Knowledgeable Community. Finally, we compare our approach with related work.

## 2 Information-centered AI and Knowledge Sharing

Traditional textbooks characterize Artificial Intelligence as a study of human intelligence and its implementation by artificial information processors. Typical approaches of AI are to isolate basic mechanisms of human information processing and design artifacts that can mimic it. This approach can be called *processing-centered AI*, as the primary focus is to understand and design a mechanism of information processing. Researchers tend to overly constrain themselves to issues of medium difficulty which are hard enough to con-

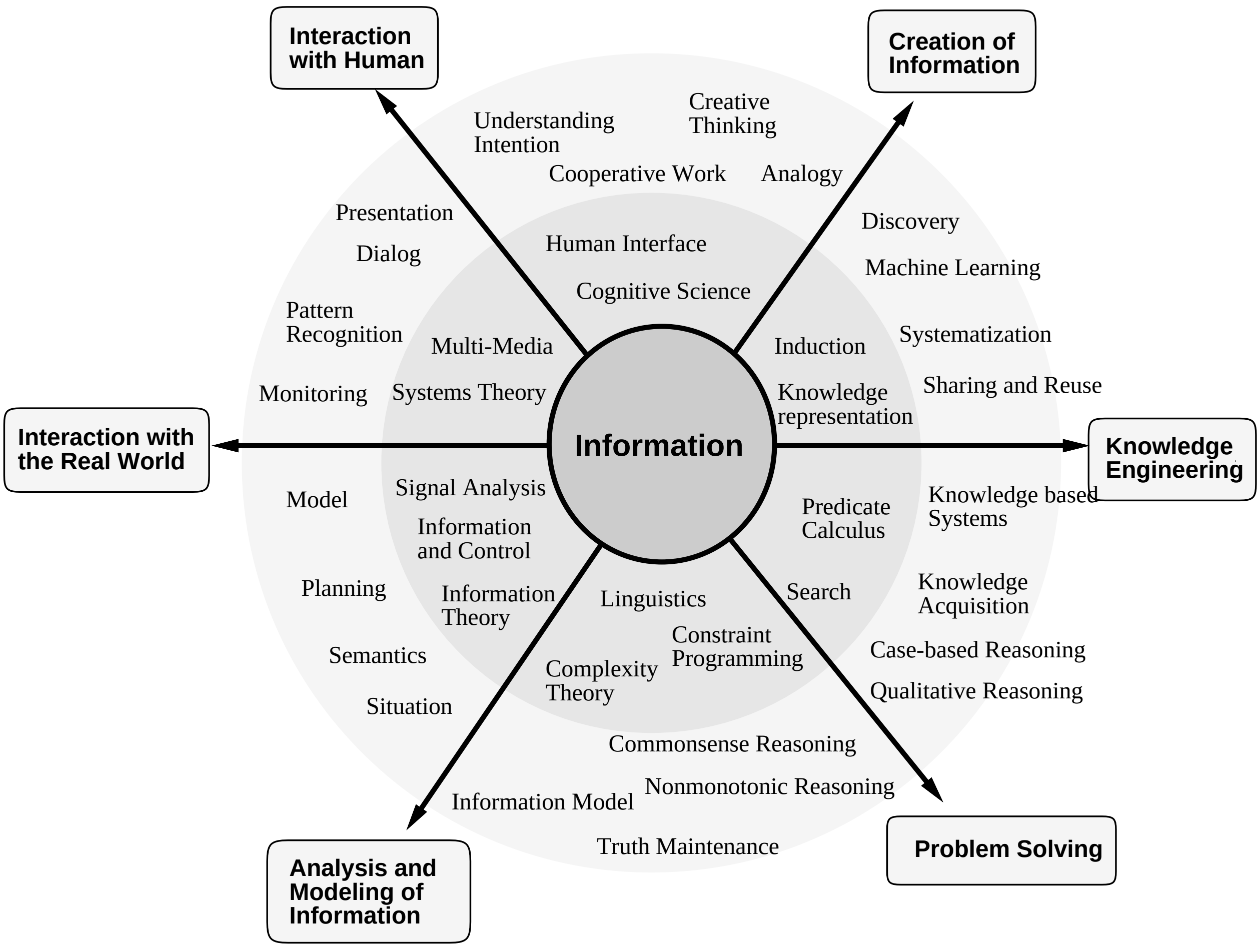


Figure 1: Information-centered AI; the central issue is information; six dimensions are defined to explore information; core issues are placed near the center of the picture.

vince people of importance of a new theory but which are not too hard to tackle.

Unfortunately, this approach runs into trouble. First, the end users cannot expect an endorsement that confirms achievement of their goal. This is because AI researchers hardly make serious efforts on developing useful systems and instead they prefer hacking sophisticated information processing mechanisms even if they are not useful. Thus, valuable information may be pushed out of systems until it is made computer-understandable. Second and more importantly, the exploitation of domain knowledge is not expected even though it is of primary importance to domain experts, for comprehensive analysis and development of knowledge base are regarded as a secondary issue and does not attract much attention. As a result, conventional AI systems often fail to keep end users whose evaluation and critical comments are invaluable to obtain insights of the problem analysis and improvement.

An alternative view, which may be called *information-centered* AI, is to think of AI as a discipline of modeling, understanding engineering, and creating the contents and structure of information, rather than information processing. Figure **??** shows research areas of information-centered AI. Analysis and modeling of information are regarded as a foundation of information-centered AI. Problem solving is concerned with theoretical issues of making use of information and deriving new information. Knowledge engineering focuses on engineering information, involving such issues as sharing and reuse of information, knowledge acquisition, organization of knowledge. Creation of information explores issues such as creative thinking or discovery. Interaction with humans handles relationship between information space and humans. Interaction with the physical world involves modeling, planning, monitoring, and understanding the physical world.

There are several advantages of information-centered AI. Unlike processing-centered AI, unstructured information only understandable to humans is regarded equally valuable and is subject to intensive study even if it is not automatically processed. The primary goal of this approach is not to fully automate the process, but to solve problems regardless of the means; it does not matter whether problems are solved automatically. Attention is focused on identifying critical problems and seeking the best way of solving them, either by full or partial computational support. Second, this view would give a better characterization of part of conventional and future AI research whose essence has been to understand the nature and structure of knowledge and information. Indeed, recent progress of research on large scale accumulation of knowledge and knowledge sharing can be captured in the context of information-centered AI, for the emphasis is on the content of knowledge rather than problem solving methods based on knowledge.

Pragmatically, even information-centered AI should accompany implementation; AI is an empirical science and AI theories have to be implemented for empirical evaluation. The way information-centered AI systems are implemented may be quite different from conventional AI systems. The purpose of implementing conventional AI systems has been mainly to test the idea. Certain aspect of intelligence is identified, isolated, and implemented as a rather independent system. In contrast, the purpose of implementing information-centered AI systems is to explore the information space. Full automatic problem solving is not always necessary. Rather, man-machine interactive problem solving is sometimes valuable either for responding immediate requirements or for gathering insights into human information processing for automation in future.

Towards this end, the following characteristics are considered to be critical to information-centered AI systems:

1. allowing heterogeneous approaches
2. being useful, even in the course of development
3. allowing many people to participate in development.

Information-centered AI can be regarded as a successor of existing technologies, in particular information filtering such information lens [?], and information resource finding [?] such as World-Wide Web [?]. Information-centered AI aims at establishing exploiting more structured and organized information space that allows intelligent information processing using AI technologies. Recent work in large scale knowledge bases [?, ?], knowledge sharing [?], and computer-aided thinking [?] are good examples in this direction.

# 3 The Framework of the Knowledgeable Community

The Knowledgeable Community is a framework of knowledge sharing and reuse based on a multi-agent architecture. Agents in the Knowledgeable Community are units of knowledge or ability of reasoning and problem-solving. Agents are spatially distributed and communicate with each other by exchanging messages. The multi-agent architecture reflects the current and future technologies of interconnected computers through high speed telecommunication networks at the hardware level and object-oriented concurrent software architecture at the software level, enabling large-scale integration of computer softwares implemented using varieties of technologies.

The Knowledgeable Community is characterized by the following features:

1. the Knowledgeable Community facilitates both information and knowledge sharing;
2. organizational structure of agents is given in which two special types of agents, facilitators and mediators, help to find agents which can provide desired service;
3. the Knowledgeable Community handles knowledge media [?] that are meaningful both to humans and to computers;
4. a testbed is given to demonstrate and evaluate the Knowledgeable Community.

## 3.1 Design Policies of the Knowledgeable Community

As participants of the Knowledgeable Community, agents are expected to cooperate with each other to solve given problems. We pose requirements on agents and their communication in order to realize such cooperation. General requirements are as follows:

**(1) Ontological relations among agents** In the Knowledgeable Community, we assume ontological relations among agents to enable sharing of knowledge, reasoning, and problem-solving. For example, *domain knowledge agents*, which manage knowledge in certain domains, are often structured hierarchically in accordance with hierarchical structure of domains.

**(2) Publication of agents' abilities** Cooperation can be achieved only if participants know each others' abilities. We request each agent to provide description of its ability. This description is used in two ways in the Knowledgeable Community. One is for agents to publish their descriptions of abilities, and the other is to reply queries about their abilities.

Furthermore, agents' abilities can be understood indirectly by ontological relations among agents.

**(3) Autonomy of agents' organization** Since there is no central organization in the Knowledgeable Community, each agent should have an ability to take care of its own actions including starting and ending. If there are ontological relations among agents, some agents instead can have abilities to take care of behaviors of related agents, such as starting, ending or even creation of agents.

**(4) Public communication languages and ontologies** In the Knowledgeable Community, various languages are used for communication, ranging from machine understandable data to natural languages, and various contents are told, varying from numerical information in engineering to literary sentences.

Communication in the Knowledgeable Community should be understandable not only by concerned agents but also by other agents to some extent so that the

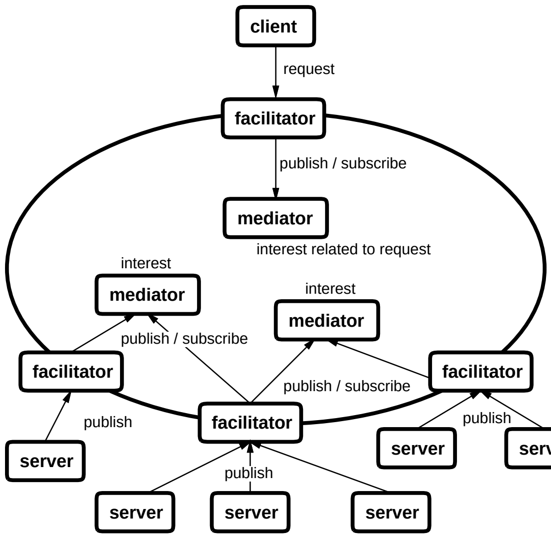


Figure 2: Organization of agents in the Knowledgeable Community

message may be interpreted by other agents and forwarded to an adequate agent whose existence has not been known to the sender.

Local languages and ontologies are allowed but should be reported to some agents in order to make messages meaningful to other agents.

**(5) Communication with unspecified addressee** The purpose of communication among agents is not communication itself, but achievement of problem solving through communication. Since agents in the Knowledgeable Community are merely sources of knowledge or reasoning, the central concern is what can be done about a problem rather than which agent can deal with it. Thus, the Knowledgeable Community should allow not only point-to-point (agent-to-agent) communiciation but communication with unspecified addressee which can reduce a priori knowledge about other agents.

**(6) Allowing various communication modes** In the Knowledgeable Community, communication varies from simple communications like question-answer pairs to long dialogues. In particular, concentrated communication may be needed to achieve cooperate work among agents. The Knowledgeable Community should provide not only sparse communication but also tightly coupled communication among agents.

### 3.2 Organization of Agents

We base the organizational structure of the Knowledgeable Community on the *federation architecture* [?, ?]. Figure ?? shows how agents in the Knowledgeable Community are organized.

We have classified agents into several types according to their roles in the agent organization (see Table ??). Roughly, the following agent types are distinguished:

- *facilitators*, like those introduced in [?, ?], which manage low level communication processing among agents.
- *mediators* which bridge between those agents requesting for a service and those providing the service.
- *ordinary agents* which provide facts about the world or derive new information by various kinds of computation or problem solving method (Table ??). Agents may be further subcategorized into *instance agents* and *class agents* depending on how they are associated with ontology.

All message transmission is made through facilitators. When the addressee is specified, the facilitator will deliver the message to the agent specified as the addressee. Otherwise, the facilitator will forward the message to mediators, which will propose the recipient or its candidates based on knowledge about other agents and ontologies.

Mediators play functions similar to facilitators in the sense that they help agents interoperate. The difference is that mediators are more like ordinary agents except for the contents of information they provide; ordinary agents provide information about the world, while mediators know what other agents do. Facilitators are more like an operating system in computer systems. They are public and handling low level communication services, and do not use heuristic knowledge. In contrast, mediators provide private information using heuristics of various kinds.

Mediators express their interests by subscribing to facilitators. On the other hand, agents publish their ability to facilitators when registered to the Knowledgeable Community. Facilitators keep and occasionally match the two kinds of publicized information.

Each ordinary agent acts either as a client or a server depending on the status of information processing. When a task is given, one of ordinary agents will serve as a client and others will be potential servers. As the process of information processing proceeds, tasks may be decomposed into subtasks and different agents may in turn serve as a client.

### 3.3 Knowledge Media as Common Knowledge Representation in the Knowledgeable Community

The Knowledgeable Community incorporates *knowledge media* which are meaningful both to humans and to computers. Some information that agents produce or receive is formally described so that computer programs can handle it, while other information only makes sense to humans who are able to interpret it. The former is inevitable for enjoying the benefit of information processing capabilities of computer software. Agent communication language, knowledge interchange format, and ontology should be shared so that computer software may interoperate. We employ KQML [?], KIF [?], and Ontolingua [?]. On the other hand, incorporating the latter type of information into systems pragmatically makes sense, for it is immedi-

Table 1: Types of agents

<table>
<tr><td colspan="3">Facilitators<br>(message handling agents)</td></tr>
<tr><td rowspan="8">(message processing agents)</td><td colspan="2">Mediators<br>(agent for messages with unspecified addressee)</td></tr>
<tr><td rowspan="7">Ordinary agents<br>(agents for messages with specific addressee)</td><td>Ontology servers</td></tr>
<tr><td>Concept associators</td></tr>
<tr><td>Information servers</td></tr>
<tr><td>Inference severs</td></tr>
<tr><td>Computation servers</td></tr>
<tr><td>Problem solvers</td></tr>
<tr><td>. . .</td></tr>
</table>

ately useful to end users and is valuable as sample data for formalization.

Essential part of an agent may be either a human or a computer software. In the former case, a human interface is provided to translate computer-oriented formal message into an informal representation so that end users can read it. In the latter, an application program interface is provided to interface application software with network of agents.

Table 2: Varieties of functions of ordinary agents

1. Information service; the following are typical functions displayed as question-answer (Q-A) or request-answer (R-A) pairs:
   (a) (Providing facts)
   Q: What is the coefficient of linear expansion of copper at 20°C?
   A: $16.5 \times 10^{-6}$.
   (b) (Providing cases)
   R: Show me typical examples of flip-flop circuits.
   A: (typical flip-flops are shown)
   (c) (Providing generic facts)
   Q: How much heat energy is radiated from object surfaces?
   A: $E = \mathrm{k}T^4$, where $E$: energy radiated, $k$: constant, $T$: the absolute temperature of the surface.
   (d) (Providing ontological information)
   Q: What are the subclasses of animals?
   A: birds, mammals, ....
2. Inference engine services such as:
   database management systems, theorem provers, language processors, belief management systems, expert system shells, learning algorithms, and so on.
3. Computation and problem solving service such as:
   numerical and symbolic computation, data analysis, diagnosis, scheduling, design, and so on; theoretically, this can be regarded as a product of an information source and an inference engine.

### 3.4 KC-Kansai as a testbed of the Knowledgeable Community

Use of good testbed is expected to highly speed up the development process. Currently, we are using travel arrangement problems as a testbed and are implementing KC-Kansai that is a multi-agent system for arranging travel in the Kansai area in Japan.

Figure ?? shows the conceptual framework of KC-Kansai. Spatially distributed agents communicate with each other through various communication media ranging from low-speed serial communication line to high-speed communication channel such as B-ISDN. Some agents provide information about a local community such as local events, local institutes, transportation, restaurants, accommodation, amusement, and so on. Others provide information processing service such as arranging meeting, travel, and so on.

This domain is interesting as a testbed, for:

1. a rich inventory of information source is involved and no single existing information source covers the information range;
2. each information source is managed individually and provides rather private information;
3. quality and characteristics of information generated by each information source are rather incoherent;
4. information sources are dynamically changing.

Several interesting applications of KC-Kansai are foreseen including travel arrangement, active yellow page, and virtual tour. Figure ?? illustrates how travel arrangement is made in KC-Kansai. In this context, KC-Kansai will help client agents in two ways. First, it will figure out a travel plan from a given specification. Second, it will monitor the execution process of a plan and deal with replanning if necessary. Personal communication media will be used for this purpose.

Active yellow page helps to find one or more service to achieve a given goal. A key question here is designing an organization of server agents so that a given set of requirements of clients can be achieved under a given set of constraints.

Virtual tour provides end users with interactive movie of local area (Figure ??). It will be made available by incorporating multi-media technique.

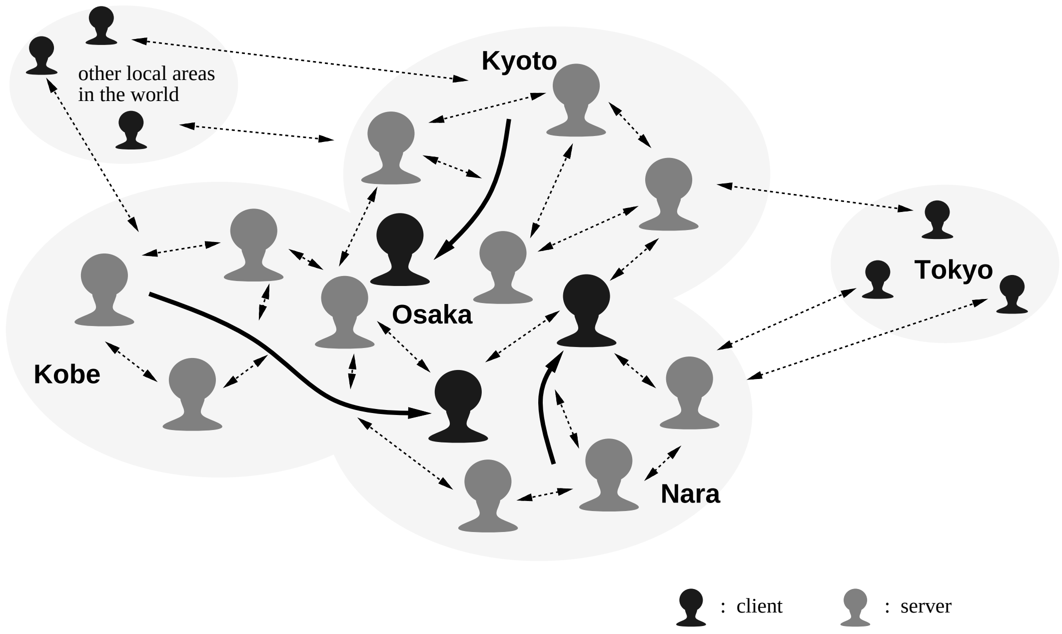


Figure 3: KC-Kansai as a testbed

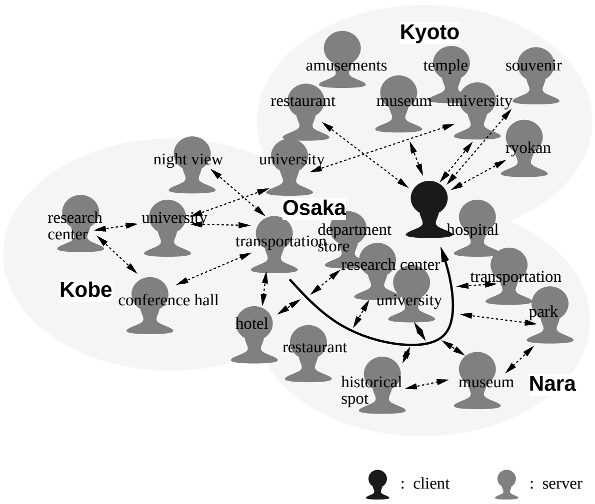


Figure 4: KC-Kansai for travel arrangement

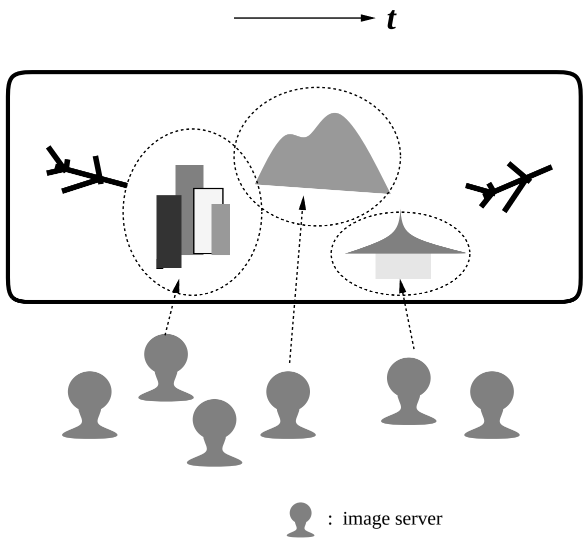


Figure 5: KC-Kansai for virtual tour

The Knowledgeable Community will be developed through multiple phases. At early stages of development, agents in the Knowledgeable Community will be mostly hand-coded.

At later stages, agents will be able to learn from others or reason about the organizational structure. Self-organization techniques might be applied to allow the organization of agents to evolve autonomously.

Currently, the development of the first stage with travel arrangement being the subject domain is in progress, which is a major focus of the rest of the paper.

## 4 The Organization of Agents for Mediation

The more agents are in the Knowledgeable Community, the harder to find one or more agent that satisfies a given goal. An adequate means should be provided for facilitating mediation between those agents requesting for a service and those providing the service. Towards this end, we propose ontology oriented agent organization and multi-stage mediation mechanism.

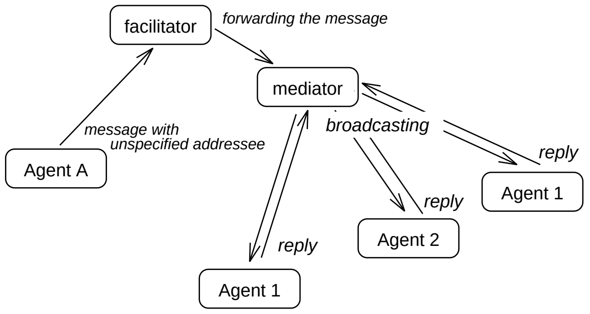


Figure 6: Mediation by broadcasting

### 4.1 Ontology-Oriented Agent Organization

Frequently, a system of concept definitions is called an *ontology*[1] [?]. An ontology plays an essential role in representing the world. We assume that agents are organized along with an ontology.

For example, the *railway* agent represents a concept which is a specialization of the concept that the *traffic* agent represent. We call the former a *sub-class agent* and the latter a *class agent*. Class agents know their sub-class agents, and sub-class agents know their class agents. There are also *instance* agents which represent individual facts in the real world. For example, the *railway* agent has some instance agents, the *railway-company-A* agent for instance.

Agents may be created not only statically but also dynamically. For example, the *travel-planner* agent can have more than one instance agent temporarily to process multiple requests.

Furthermore, there are special agents called *ontology servers* which maintain all relations in an ontology. Not all concepts in an ontology are associated with an agent. An ontology server provides information about concepts and relations.

### 4.2 Multi-stage Mediation

The Knowledgeable Community allows agents to send a message without specifying the addressee. Such messages are sent to special agents called *mediators*, which decide which agent is appropriate to receive them and forward the messages to it.

#### Mediation by broadcasting

The simplest mediation is to ask all or some agents whether they can respond the message. A mediator collects replies and decides appropriate agents to send the message (see Figure **??**).

#### Mediation by class hierarchy

If there is a class hierarchy of agents, mediators forward the message to some class agent. Since a class agent knows abilities of its sub-class agents, it can

[1] There may be more than one ontology. Different ontologies represent different characterizations of the world.

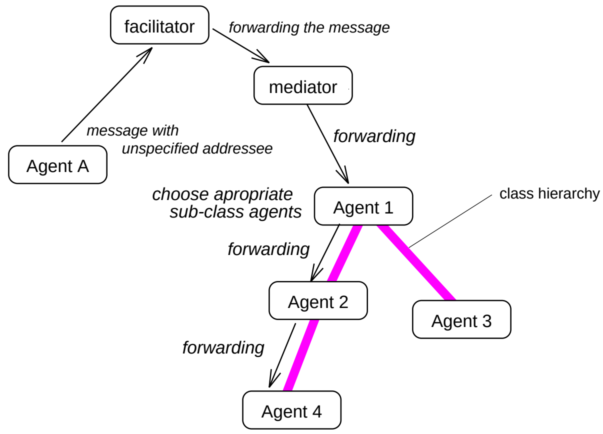


Figure 7: Mediation by class hierarchy

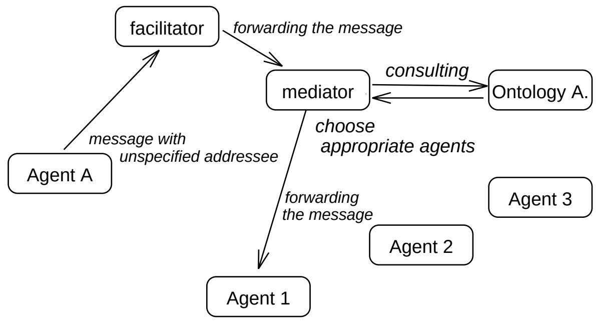


Figure 8: Mediation with ontology server

choose appropriate sub-class agents and forward the message to them (see Figure **??**). If there are more than one layer in the concept hierarchy, the message can be passed through multiple class agents.

#### Mediation by ontology

The role of a mediator is to find appropriate agents related to the given message, by asking relations among agents and those among concepts to an ontology server (see Figure **??**). If there exists an agent associated with the concept in the message, an ontology server returns the name of the agent. Otherwise, it will seek agents which may be related to the concepts in the given message.

#### Mediation by heuristics

A rigorous class hierarchy is not powerful enough to cope with vaguely specified requests. Such requests will be handled by mediation by heuristics that calls for such agents as *concept associators* or *case-based mediation*. Concept associators use concept spaces which give distances between concepts based on various criterion, allowing to retrieve concepts relevant to a given set of concepts. Case-based mediation is a mechanism of reusing previous experiences of mediation in which

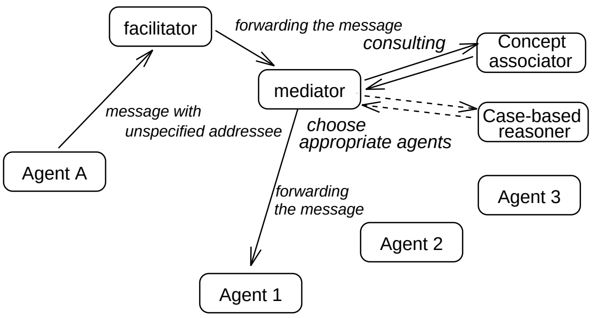


Figure 9: Mediation by heuristics

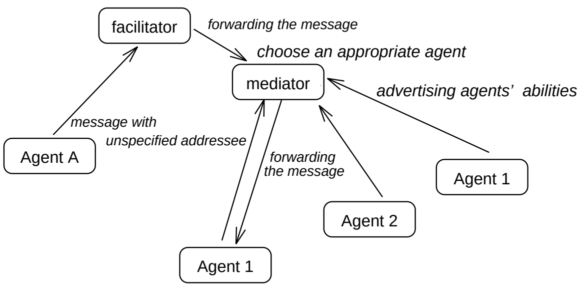


Figure 10: Mediation by descriptions of agents' abilities

nontrivial reasoning is performed, either automatically or manually, either to find the target. (Figure **??**).

**Mediation by descriptions of agents' abilities**

A mediator interprets descriptions of agents' abilities to decide which agents are appropriate to the message (see Figure **??**). This type of mediation can be combined with mediation by heuristics to process collected descriptions of agents' abilities.

### 4.3 Examples of communication and mediation

Consider a system of agents shown in KC-Kansai **??**. We show how varieties of messages can be handled below. Although there may be more than one facilitator or mediator, we simply describe as if there there were only one facilitator and mediator, respectively.

(1) "Tell me your phone number" ⇒ the *Hotel-A* agent

This is almost the simplest case, for the addressee is explicitly given and the facilitator will deliver the message directly to the *Hotel-A* agent, an agent associated with *Hotel-A*.

(2) "Tell me the phone number of Hotel-A" ⇒ ?

To handle this message, the mediator asks the ontology server for how features of *Hotle-A* can be accessed. In this case, since *Hotel-A* is realized as the *Hotel-A*

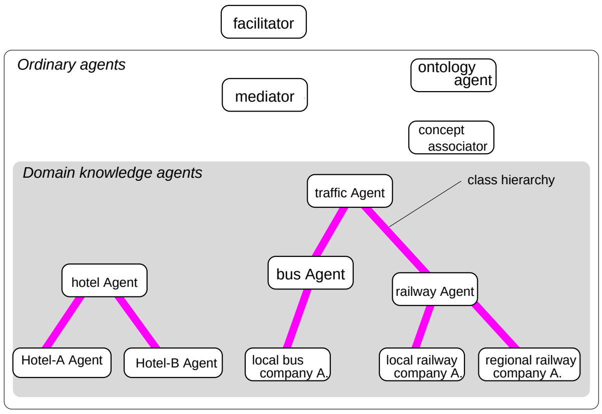


Figure 11: Examples of KC-Kansai agents

agent, the ontology server replies with the name of the agent, allowing the mediator to deliver the message to the *Hotel-A* agent.

(3) "Tell me hotels where room price is within 10,000 Yen" ⇒ ?

The mediator forwards this message to the *hotel* agent, a class agent associated with the class *hotel*. Then, the *hotel* agent broadcasts the message to its instance agents each of which represents an individual hotel, and collects the responses that meet the given constraint.

(4) "Tell me how long it will take from Station-A to Station-B" ⇒ ?

We have assumed that *Station-A* and *Station-B* are sub-classes of class *station*, but neither of the three is associated with an instance agent. In this case, the ontology server will seek indirect ways to obtain the information. For example, the ontology agent may know that information about stations is accessible from the *railway* agent. Then, the mediator will create a new message, send it to the *railway* agent, and transform the response so that it may be adequate as a response to the original message.

(5) "Propose a one-day trip which involves discussing about research and education, appreciating Japanese culture, and some refreshing time" ⇒ ?

Since this is a vaguely specified request as far as visiting locations are concerned, mediation by heuristics will be needed by the *travel planner* agent to figure out a set of locations to visit. In the context of KC-Kansai, relevant concept spaces are such as those given in Figure **??**. A concept associator may be invoked to find a minimal cluster of concepts that may maximally fit a given set of key words: "discussing about research and education", "appreciating Japanese culture", and "refreshing". In this case, the concept associator may find that pairs of visiting locations, NAIST and Nara

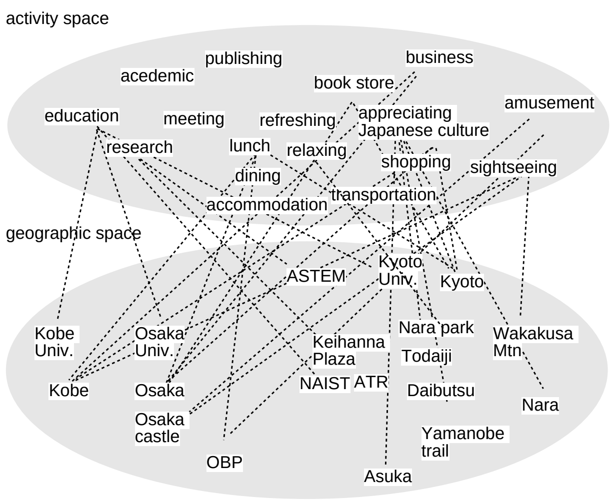


Figure 12: Concept spaces relevant to KC-Kansai

park, best cover a given request, as follows:

| | | | | |
|---|---|---|---|---|
| discussing about research and education<br><br>appreciating Japanese culture<br><br>refreshing | ~ | visiting university<br><br>appreciating Japanese culture<br><br>relaxing | ~ | NAIST<br><br>Nara park |

and returns them as a response. Now the *travel planner* agent comes to have a rather concrete plan as far as the visiting locations are concerned and will elaborate the plan using less heuristic mediation methods.

## 5 Building up the Knowledgeable Community

Development of the Knowledgeable Community is a long-term goal and is expected to take a long period. In order to keep the Knowledgeable Community useful even at intermediate stages, we develop agents in the Knowledgeable Community in a bottom-up, inductive manner. It is our belief that building up the concept structure is a long-term trial-and-error process based on rich data, rather than top-down, deductive process [?].

Agents at the first stage only provide unstructured information such as natural language texts, voices, or images which may only make sense to humans. These kinds of information is easy to obtain and useful to humans, even though agents cannot understand the contents.

As the stage of development proceeds, the more sophisticated techniques will be used to improve agents so that they can respond with more structured information, allowing new agents (or *agencies* [?]) to be implemented which may derive new information from structured information.

Currently, the building up process is manually performed. We could use induction algorithms to induce generic rules from data. This is still semi-manual, for evaluation and incorporation of hypotheses are left for humans. In future, this process might be further exploited by introducing some kind of self-organization facilities.

## 6 Related Work

The Knowledgeable Community can be regarded as a successor of several technologies.

Compared with multi-media technology, the Knowledgeable Community is distinguished from existing technologies by the degree information and knowledge sources are structured and organized. The current multi-media technology assumes that humans play an essential role in interpreting and manipulating information. Although this approach brings about direct implications to the end users, it does not decrease intellectual information processing loads. In contrast, the Knowledgeable Community aims at knowledge media that are understandable and manipulatable both by humans and computers. AI techniques will be used to structuring and organizing information so that computer software may handle it, decreasing the various load required to make use of potential information.

The agent-oriented technology is similar to the object-oriented technology. However, there are two differences. First, agents are active in the sense that they understand and react the environments autonomously, while objects are passive entities that only respond incoming messages. Second, instance objects are almost empty, just a set of individual values in object-oriented systems, while instance objects are crucial in the agent-oriented technology. In the agent-oriented technology, instance agents are not simply a set of individual values but also individual entities that may persist and evolve over time.

We have chosen the platforms developed by the DARPA knowledge sharing and reuse initiative as a base. The net progress in the Knowledgeable Community is the the use of knowledge media and powerful mediation mechanism.

## 7 Conclusion

In this paper, we have proposed information-centered AI as opposed to processing-oriented AI, and presented a framework of the Knowledgeable Community. We have focused on the organization of agents for mediation and described two methods: ontology-oriented agent organization and powerful mediation mechanism.

## References

[1] Tim Berners-Lee, Robert Cailliau, Jean-Francois Groff, and Bernd Pollermann. World-wide web: The information universe. *Electronic Networking: Research, Applications and Policy*, 1(2), 1992.

[2] Mark R. Cutkosky, Robert S. Engelmore, Richard E. Fikes, Michael R. Genesereth, Thoas R. Gruber, William S. Mark, Jay M. Tenenbaum, and Jay C. Weber. PACT: An experiment in integrating concurrent engineering systems. *IEEE Computer*, January 1993:28–38, 1993.

[3] Peter B. Danzig, Katia Obraczka, and Shih-Hao Li. Internet resource discovery services. *IEEE Computer*, September 1993:8–22, 1993.

[4] T. Finin, J. Weber, G. Wiederhold, M. Genesereth, R. Fritzson, D. McKay, J. McGuire, P. Pelavin, S. Shapiro, and C. Beck. Specification of the KQML agent-communication language. Technical Report EIT TR 92-04, Enterprise Integration Technologies, 1992. (Updated July 1993).

[5] Michael R. Genesereth and Richard Fikes. Knowledge Interchange Format version 3.0 reference manual. Technical Report Logic-90-4, Computer Science Department, Stanford University, 1990.

[6] T. R. Gruber. Ontolingua: A mechanism to support portable ontologies. Technical Report KSL 91-66, Stanford University, Knowledge Systems Laboratory, 1992.

[7] Thomas R. Gruber. Ontolingua: A mechanism to support portable ontologies Version 3.0. Technical report, Knowledge Systems Laboratory, Stanford University, 1992.

[8] R. V. Guha and D. B. Lenat. Cyc: A midterm report. *AI magazine*, pages 32–59, Fall 1990.

[9] Koichi Hori. A system for aiding creative concept formation. *IEEE Transactions on Systems, Man, and Cybernetics*, 24(5), 1994. (to appear).

[10] Douglas B. Lenat and R. V. Guha. *Building Large Knowledge-based Systems*. Addison-Wesley, 1989.

[11] Thomas W. Malone, Kenneth R. Grant, Franklyn A. Turbak, Stephen A. Brost, and Michael D. Cohen. Intelligent information-sharing systems. *Communications of ACM*, 30(5):390–402, 1987.

[12] Marvin Minsky. *Society of Minds*. Simon & Schuster, Inc., 1985.

[13] Robert Neches, Richard Fikes, Tim Finin, Thomas Gruber, Ramesh Patil, Ted Senator, and William R. Swartout. Enabling technology for knowledge sharing. *AI Magazine*, 12(3):36–56, 1991.

[14] R. S. Patil, R. E. Fikes, P. F. Patel-Schneider, D. McKay, T. Finin, T. R. Gruber, and R. Neches. The DARPA knowledge sharing effort: Progress report. In Charles Rich, Bernhard Nebel, and William Swartout, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the Third International Conference*. Morgan Kaufmann, 1992.

[15] Mark Stefik. The next knowledge medium. *AI Magazine*, 7(1):34–46, 1986.