

Representations of Health Concepts: A Cognitive Perspective

Jiajie Zhang, Ph.D.

Associate Professor, School of Health Information Sciences
University of Texas Health Science Center at Houston
7000 Fannin, Suite 600
Houston, TX 77030
Tel: (713) 500-3922; Fax: (713) 500-3929
Email: Jiajie.Zhang@uth.tmc.edu

Abstract

Objective: This paper studies the differences between controlled medical vocabularies that are designed as external artifacts and the mental concepts that are inside users' heads and used by users for reasoning, decision making, diagnosis, and treatment.

Design: The major theories of concept representations developed in cognitive science were reviewed, analyzed, and compared with the major controlled medical vocabularies developed in medicine.

Results: It was found that there are significant mismatches between controlled medical vocabularies that are designed as external artifacts and the mental concepts that are inside users' heads and used by users for reasoning, decision making, diagnosis, and treatment.

Conclusions: Controlled medical vocabularies should be designed with systematic considerations of the cognitive structures and processes of the users. Without such considerations, the designed vocabularies will not be appropriate for people because they are hard to use, although they may or may not be appropriate for machine processing.

KEYWORDS: *Controlled medical vocabularies, terminology, knowledge representation, health concept, cognitive science*

Introduction

Electronic medical records (EMR) have the potential to make a highly significant contribution to the advancement of medicine and to the improvement of the quality of health-care. An ideal EMR would provide complete, accurate, and timely data, alerts, reminders, clinical decision supports, medical knowledge, communications, and other aids at all points of care for all healthcare professionals at all times in a way the quality of healthcare can be dramatically improved. In order to achieve these promised functions of EMR, one necessary condition is a structured medical vocabulary that serves as the foundation of EMR. Many classification systems have been proposed and implemented [1, 2, 3, 4]. However, no current system can capture the full scope of medical knowledge. To address the potential problems in the rapid growth of vocabulary contents, Cimino [5] proposed a set of desiderata for the next generation of standard, reusable, multipurpose controlled vocabularies. On the scientific foundation side, Chute and colleagues [6] argue that if we want to achieve reliable outcomes and efficient assessment of data, we need to pay significant attention to the basic science of representing what we do to patients. Coiera [1] made similar arguments on the important roles of the cognitive foundation of medical concept representation. Along this line, Patel and Kushniruk [7] examined the cognitive issues of how users understand, navigate, and

communicate medical knowledge. In a recent study published in this issue, Patel, Arocha, & Kushniruk [8] showed that knowledge representations of medical problems by physicians and patients were radically different. Such a mismatch could have an impact on the nature of medical decision strategies used for patient care.




This paper examines the representational issues of health concepts from a cognitive perspective. It starts with a definition of representations. Based on this definition, the differences between two types of concepts and vocabularies are described. Then seven cognitive theories of concepts and their implications for the representations of health concepts are discussed. Finally, the implications of expert-novice differences in concept representations are discussed.

Representations

A representation is something that stands for something else. It is a mapping between a represented world (that which is to be represented) and a representing world (that which does the representing). A representation must specify which aspects (objects, properties, and relations) of the represented world

are to be modeled in the representing world, as well as how the representing world carries out this mapping [9] Table 1 shows a simple example of representations. In the represented world, there are two objects, two properties for each object, and two relations between the two objects. These two objects and their properties and relations can be mapped to the representing world in different ways. In Representing World 1 they are represented by words and numeric values, whereas in Representing World 2 they are represented by pictures. In Representing World 1, more knowledge is stored in internal representations (memory in the head). For example, it is necessary to retrieve numerical facts from memory to compare the heights represented by numbers. In Representing World 2, more knowledge is available in external representations (pictures). For example, the weights can be compared by visually inspecting the sizes of the two circles. Under many circumstances Representing World 2 is more efficient than Representing World 1 because less mental effort, though always true, typically leads to higher efficiency in task performance.

Table 1. Representation

Represented World		Representing World 1	Representing World 2
Objects	Object 1, Object 2	John, Eric	
Properties	Height	70 inches, 65 inches	
	Weight	180 LB, 140 LB	
Relations	TallerThan(1, 2)	Relative numeric values	Relative lengths of bars
	HeavierThan(1, 2)	Relative numeric values	Relative sizes of circles

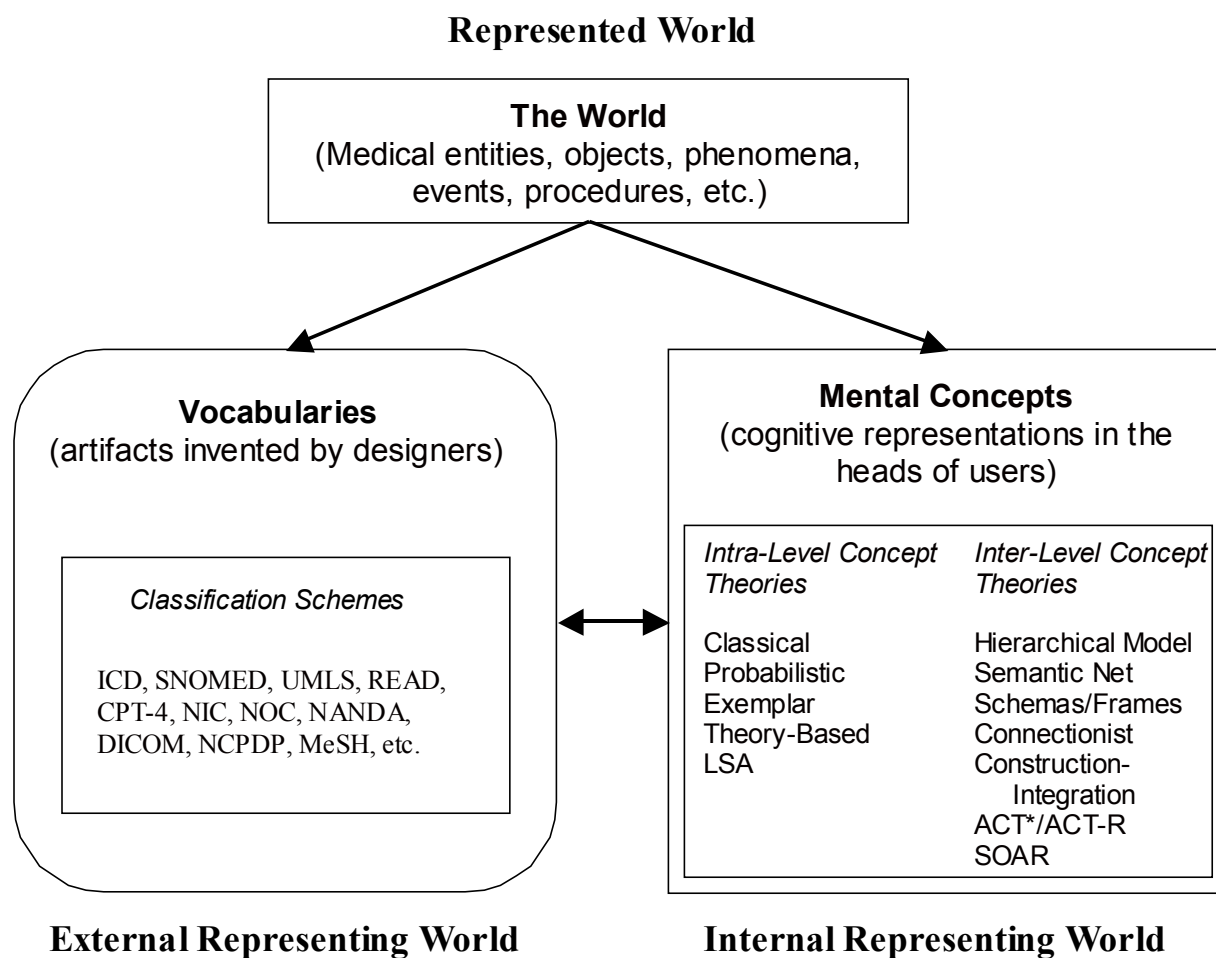


Figure 1. The relations between represented and representing worlds. See text for details.

Figure 1 shows the relations between represented and representing worlds for the complex medical domain. The represented world is the world of all medical entities, objects, phenomena, events, procedures, etc. They are the things to be represented. The representing world can be the vocabularies invented by designers of classification systems. In this case, the things in the representing world are external to the human mind. They are artifacts created by human beings to carve up and categorize the medical domain in a systematical way. The representing world can also be the mental concepts in the minds of users. In this case, the things in the representing world are internal

to the mind. They are the concepts acquired by users through learning and memorization, and they are the concepts that users use to perform diagnoses, reasoning, decision making, problem solving, treatment, and so on. The external representations of vocabularies and the internal representations of mental concepts have different properties that may affect the processing of information in a non-trivial way [10, 11, 12].

The critical issue is the relation between the two representing worlds: vocabularies in external classification systems and mental concepts inside the head of users and used by users [7, 13]. Let us consider the following four cases.

In the first case, external vocabularies in classification systems are identical to the internal mental concepts in users' minds, and the external vocabularies are copied from the internal mental concepts. In this case, the mental concepts are optimized for human processing. However, the external vocabularies are unlikely to be appropriate for machine processing (e.g., in EMR) because the regularities of external vocabularies reflect the structures and constraints of mental representations that are usually incompatible with machine representations. One example of this is the multiple ways of coding the same meaning by human users in a machine system (e.g., SNOMED International) that cannot easily figure out that the multiple representations are equivalent in meaning. One example described by Hammond & Cimino [14] is the representation of "acute appendicitis", which can be coded as a single disease term, as a combination of a modifier (acute) and a disease term (appendicitis), or as a combination of a modifier (acute), a morphology term (inflammation) and a topography term (vermiform appendix). All of these representations are correct, but there is no formal way in SNOMED to know that their meanings are equivalent.

In the second case, the external vocabularies are identical to the internal mental concepts, but the internal mental concepts are copied from the external vocabularies. In this case, the regularities of internal mental concepts reflect the structures and constraints of external vocabularies. The external vocabularies may be optimized and appropriate for machine processing. However, it is not likely that they are also appropriate for human processing, unless the structures of the human mind are systematically considered during the design of the external vocabularies. ICD-9 (International Classification, of Diseases, Ninth Edition) developed by the World Health Organization [15] is

primarily a vocabulary for the classification of diseases for statistical purposes. It is based on a strict hierarchy representation, which does not match well with the structure of patient information processed by human users in practice. If it is used as a terminology system for the coding of patient information, its artificial structure has to be learned and memorized by human users.

In the third case, the external vocabularies are different from the internal mental concepts, and they have evolved independently. In this case, if the external vocabularies are designed to optimize machine processing and the internal mental concepts have evolved to optimize mental processing, the success of the two combined systems will depend on whether there is an efficient translation mechanism between the external vocabularies and the internal mental concepts. If such a mechanism does not exist, it would be very difficult for machines to process human mental concepts and for human beings to process machine vocabularies, as in the case of natural languages vs. programming languages. Probably due to its obvious problem, none of the current medical vocabulary systems fall into this category.

In the fourth case, the external vocabularies overlap with the internal mental concepts, and they have evolved interactively. This is a realistic situation and it more or less reflects the status of current classification systems. For example, users acquire mental concepts through internalization of external vocabularies (e.g., through textbooks or handbooks), which in turn might be designed with some consideration of the cognitive properties of the mind. The problem with this situation is that the external vocabularies and the internal mental concepts are neither optimized for machines nor for human minds. Most of the major systems in use today fall into this category (e.g., ICD-9, Read, UMLS, SNOMED-RT, etc.).

Although it may not be impossible to optimize a vocabulary for both machines and human minds, it would be a long way to go because neither a vocabulary for machines nor for human minds has been optimized independently yet. Coiera [1] even argued that creating a task-free or universal vocabulary for every thing is doomed to fail because terms are subjective, context-dependent, purposive, and adaptive through evolution.

In the design of classification systems of medical vocabularies, we need to be careful about the structures of the systems. It has been shown that there is a strong cognitive phenomenon called representational determinism [10, 16], which is that the format and structure of a representation can guide, constrain, and to some extent determine the way the mind functions. If the structure of a classification system matches the user's cognitive properties, it can enhance the performance of the user. But if there is a mismatch, the user's performance could be greatly hindered. As shown in Figure 1, many classification schemes have been developed for external medical vocabularies, and many cognitive theories have also been developed for mental concepts. Currently there is a lack of communication between medical informatics practitioners who develop classification schemes and cognitive scientists who develop cognitive theories. The rest of this paper briefly reviews the fundamental phenomenon and several cognitive theories of mental concepts that could be potentially considered by the designers of medical vocabularies.

Mental Concepts

Mental concepts are important components of human intelligence. A mental concept is the mental representation of a category, which is a set of entities. The process of forming a mental concept is called categorization. Without categorization, understanding, reasoning, prediction, language, and

other high level intelligent behaviors are impossible. Even in concrete thinking situations, perceptual categorization based on perceptual learning and frequency adaptation is still needed for intelligent behaviors. If every apple is a different apple that is new to a creature, this creature cannot exhibit intelligent behavior such as prediction and expectation. In particular, concepts have the following functions. First, concepts provide classifications, that is, we can decide whether an instance belongs to a concept. Second, concepts provide understanding and explanation. Once we know that an instance belongs to a concept, we will know the properties of the instance. For example, once we know that an object is an apple, we know it is edible. Third, concepts support reasoning. Given that Sam is a dog, we can answer the question "is Sam a mammal?" The answer to this question can be inferred from the relations between concepts. Fourth, concepts provide organizations. Concepts are artifacts that artificially carve the nature into organized chunks. For example, we use limited number of color names to divide the continuous color spectrum. Fifth, concepts support communication. With the same representations of concepts, people are able to communicate with each other and learn indirect experience through communication.

Many theories have been developed for mental concepts [17, 18, 19, 20, 21, 22, 23, 24]. They basically fall into two categories (see Figure 1). The first category includes theories for intra-level concept representations. They are mainly concerned with the internal representational structures of individual concepts. Examples include the classical view, the probabilistic view, the exemplar view, the theory-based view, and the Latent Semantic Analysis (LSA) approach. The second category includes theories for inter-level concept representations. They are mainly concerned with the relations among concepts. Examples include the

hierarchical model, semantic net, frames and schemas, connectionist models, construction-integration theory, ACT* and ACT-R, and SOAR.

Intra-Level Concept Theories

This subsection briefly describes four intra-level concept theories that are most relevant to medical concepts.

Classical View. According to the classical view, a concept is defined by a set of necessary and sufficient features. For example, triangle is defined by 3 sides and a sum of 180 degrees for the interior angles. Once we know that an object is a triangle, then we know that it has three sides and the sum of its interior angles is 180 degrees (necessary features). On the other hand, if we know that an object has three sides and its interior angles add to 180 degrees (sufficient features), then we know that this object must be a triangle. The classical theory of concept representation is based on definitions of concepts that defined by necessary and sufficient features, which are mathematically elegant because they can be treated mathematically in a systematical way. One good aspect of this theory is that features of concepts are nested in subset relations. For example, features of bird are nested in those of robin because robin is a subset of bird. Another good property is that a concept is a representation of an entire class, not a set of exemplars. However, this theory is too restrictive to be a general theory. This theory does not allow disjunctive concepts. For example, concept [(red, square) or (blue, diamond)] has no defining features. And concept [(red, square, large) or (blue, square, small)] has no jointly sufficient features. Examples of disjunctive concepts include chair, furniture, game, and many other concepts that are usually described as family resemblance [25]. The principle of family resemblance is that members of a category tend to share properties with each other but there is no set of properties that each and

every member has to have. In other words, different members of the category tend to share different properties. The role of the classical view for medical concepts is very limited because most medical concepts are not well structured and many of them cannot be represented by necessary and sufficient features.

Probabilistic view. According to the probabilistic view [26], features of a concept are salient ones that have a substantial probability of occurring in instances of the concept. If x has some critical sum of weighted features of y , then x is a y . In addition, the representation of a concept is an abstraction process in which a prototype is formed through a dynamic process, not static descriptions of features. For example, every feature of a concept has a weight, whose value (between 0 and 1) can dynamically change in response to the frequencies of encounters of instances of the concept. One important point of this view is the notion of prototypes [25]. A prototype is the best example of a category; it possesses all characteristic features of a category; and it is the central tendency and average of the concept. In this view, a concept is organized around its prototype in terms of a set of features with appropriate weights that are representative of the prototypes of the concept. This view can easily explain the typicality effect that a more prototypical instance can be processed more efficiently [27, 28]. For example, people are faster to answer the question “is robin a bird?” than “is chicken a bird?” because robin is a more typical bird than chicken. However, this view cannot easily explain context-dependent concepts. For example, a harmonica is a typical musical instrument in the context of a campfire but not a typical musical instrument in a concert.

Exemplar view. According to the exemplar view [29, 30], the representation of a concept consists of separate descriptions of

some of its exemplars (either instances or subsets). In this view, classification is based on similarity to a particular exemplar. People initially learn some examples of a concept and then classify a new example based on how similar it is to the already learned particular examples of the concept. That is, a new instance elicits similar old examples and it is assumed that similar instances belong to the same category. For example, you might classify one diagnosis as flu because it reminds you of a case that you know is flu. As another example, the knowledge that large birds are less likely to sing than small birds may be derived from exemplars of small and large birds. The exemplar view is more successful than the probabilistic view, partially because it is more conservative with respect to discarding potentially relevant information (i.e., keeping detailed information without too much abstraction). However, it also has problems. For example, it cannot explain how concepts are created in the first place. That is, it does not have a mechanism for the formation and creation of concepts [24, 26].

Theory-based view. The theory-based view was developed in the studies of how concepts are used in reasoning [18, 20, 23]. In this view, the organization of concepts is knowledge-based and theory-driven, and categorization is an inference process, not a similarity judgment. For example, “children, money, photo albums, and pets” belong to a concept, which is “things to take out of one’s house in case of a fire”. In this example, classification is not based on a direct matching of properties of the concept with those in the example, but rather requires that the example have the right explanatory relationship to the theory organizing the concept. One good feature of the theory-based view is that it addresses the question of why we have the concepts we have because a theory-based concept is created by explanatory relations among the instances of the

concept. In addition, it provides a natural way in which concepts may change, that is, through the addition of new knowledge and theoretical principles.

Inter-Level Concept Theories

There are many theories for inter-level concepts. Only three are described here.

Hierarchical Model. According to the hierarchical model (see Figure 2), concepts are represented by nodes and links in a hierarchical structure. Each node is an atom concept and the link between two nodes is can be a subset relation called *is-a*, part-whole relation called *part-of*, belonging relation called *has-a*, or other relations. For example, robin *is-a* bird and bird *is-an* animal. Here we have three concepts at three levels linked by two *is-a* relations. The hierarchical model is basically a taxonomy of things. Although it captures some regularities of mental concepts, it is not a very accurate model of how concepts are represented in people’s mind. It cannot explain the typicality effect. For example, people are faster to answer the question “is robin a bird?” than “is chicken a bird?” because robin is a more typical bird than chicken. The hierarchical model would predict that there is no difference in response times for these two questions because both the robin and bird pair and the chicken and bird pair have one *is-a* relation. The hierarchical model cannot explain the reversal effect, either. For example, people are faster to answer the question “is chicken an animal?” than “is chicken a bird?” even if the former question spans three levels with two *is-a* relations whereas the latter one only spans two levels with one *is-a* relation. The third effect that cannot be explained by the hierarchical model is the relatedness effect. For example, people are faster to answer the question “is bear a bird?” than “is bat a bird?” even if both questions span three levels with two *is-a* relations. Descriptions of these three effects and a number of research studies can be found in [26].

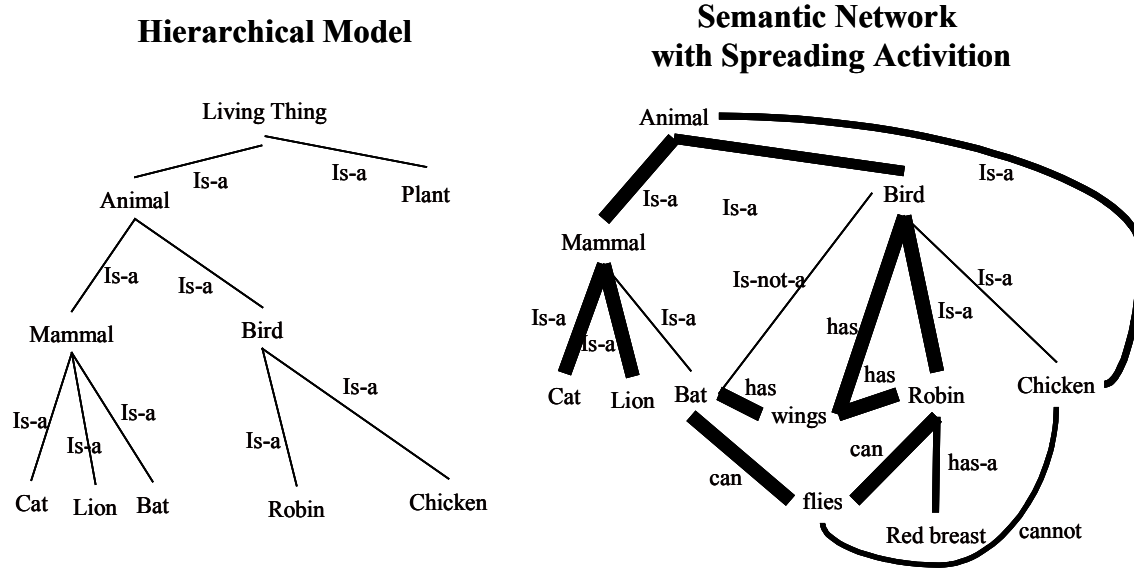


Figure 2. Hierarchical Model and Semantic Network Model with Spreading Activation

One related issue to the hierarchical model is the notion of basic level concepts. In a concept hierarchy, concepts are organized hierarchically: higher-level categories subsume lower-level categories, and different levels of concepts have different degrees of abstraction. Figure 3 shows a hierarchical structure. One of the levels in Figure 3 is called basic level, which is cognitively privileged. Basic level concepts have the following properties [31]. (1) They maximize information: superordinate level concepts have too little information and subordinate level concepts do not provide much more information. (2) They are neither the most abstract nor the most specific, but intermediate. (3) They are the first to be learned. (4) They are the objects that are naturally named. (5) They are consistent across cultures. (6) The basic level is the highest level in which the instances all share the same parts, overall shape, and associated motor movements. One obvious application of basic level concepts is to design classification systems of medical concepts such that the organization, learning, retrieval, and display of medical knowledge are centered on basic

level concepts. Coiera [1] made similar arguments with several medical examples.




Figure 3. Basic level concepts.

Semantic Network with Spreading Activation. The semantic network model with spreading activation (see Figure 2) is an extension of the hierarchical model described above. Unlike the strict hierarchical model, it has connections across more than two levels. In addition, its links have weights, with larger weights indicating more active links thus making the response times shorter. With these two added features, the semantic network model can explain the typicality effect, the reversal effect, and the related-

ness effect that cannot be explained by the hierarchical model. The typicality effect can be explained by the stronger link (larger weight) between bird and robin than between bird and chicken; the reversal effect can be explained by the direct link between animal and bird; and the relatedness effect can be explained by the strong links between bat and the features of bird (has-wings and can-flies). However, one weakness of this model is that the weights of the links cannot be automatically and systematically set.

Connectionist view. According to the connectionist view, concepts are not represented in discrete, localist units such as features and instances [32, 33]. Rather, concepts are represented as activation patterns in parallel and distributed networks of neuron-like units that connect to each other and excite or inhibit each other through dynamic processes. This view is closer to the underlying brain mechanisms of concept processing than all other views or models because its structures and mechanisms are inspired by the structures and mechanisms of the neural networks in the brain. In this view, concept is not just a structure embedded in the connection strengths between units; it is also the dynamic process of the network [22]. Concepts are no longer binary: an instance can belong to a concept to some degree, and an instance can belong to different concepts in different contexts. Concepts are no longer static: they can be modified dynamically, they are adaptive, and the acquisition process of concepts is naturally accounted for by the learning mechanisms of connectionist networks. The connectionist view has the potential to become the dominant psychological theory of concepts because it has better explanation and prediction power. However, in implementations, one big hurdle is that it is hard to build a connectionist system of concepts that can easily scale up. In a typical training of a connectionist network, the training set is usually problem specific.

For a real world problem, however, there are typically numerous variations of a problem.

Expertise and Concepts

Another cognitive issue of mental concept representation is the different ways of representing and retrieving concepts by experts and novices. A classic study by [34] shows that novices used surface features (e.g., physical shapes) to categorize problems while experts used underlying principles (e.g., physics laws) to categorize problems. Another study by Luria [35] showed that people with formal schooling used abstract information to categorize objects. For example, given hammer, log, and saw, these people grouped hammer and saw together because they were both tools. However, people with little formal schooling used practical thinking to categorize objects: they grouped saw and log together because saw was used to cut log. One implication of these studies is that the design of classification systems of medical concepts should consider the different thinking and reasoning styles of expert physicians and less experienced medical students and residents, and between physicians and patients. This point is made and demonstrated clearly by Patel, Arocha, & Kushniruk [8], who showed that knowledge representations of medical problems by physicians and patients were radically different, which could potentially affect the communication between physicians and patients and the way medicine is practiced.

Discussion and Conclusion

This paper discussed the differences between medical vocabularies that are designed as external artifacts and the mental concepts that are inside users' heads. It argues that in order for controlled vocabularies to be more usable for people, they should be designed with systematical considerations of the cognitive structures and processes of the users. Without such considerations, the de-

signed vocabularies will not be appropriate for people to use, although they may or may not be appropriate for machine processing.

This paper described four views of mental representations of intra-level concepts and three views of mental representations of inter-level concepts. Controlled medical vocabularies currently being developed cover not just simple singular concepts but also compound concepts and full range of medical knowledge. Although the cognitive theories are mainly for simple concepts, they still have values in the design of complex vocabularies. Other cognitive issues such as basic-level concepts and expertise were also discussed. It is argued that the properties of basic-level concepts and the expert-novice differences in concept representations should be incorporated to the design of controlled medical vocabularies.

Medical terminology is an issue of knowledge representation in a complex domain. Knowledge representation has been a fundamental issue in all sub-areas of cognitive science, including cognitive psychology, artificial intelligence, linguistics, cognitive neuroscience, anthropology, sociology, philosophy, and cognitive engineering. Most of the empirical studies on knowledge representation in cognitive science are basic science research that uses simple tasks performed by either college undergraduate students or everyday people. This can be clearly seen from the cognitive studies of concepts reviewed in this paper. Although expert systems developed in artificial intelligence research are for complex domains such as medicine [36, 37], they are typically not directly supported by empirical studies. The development of medical terminologies is facing the same problem. As stated earlier, without a deep understanding of how medical concepts are created, used, and represented by people and without strong and systematical empirical backing, medical terminology systems cannot go very far.

This problem is especially critical to medical terminology systems that are considered for EMR (e.g., SNOMED and READ). Clinicians who use EMR will not likely use coding systems that do not match their clinical language and are difficult to use. Without first hand, prompt capture and coding of data into EMR by clinicians, EMR will not be able to fulfill many of its promised functions, as described at the beginning of this paper. In order to develop a cognitively oriented medical terminology system, knowledge and expertise are needed both in medicine and cognitive science. This could be potentially carried out by double experts who are specialized in both medicine and cognitive science. Such double experts can be trained in a medical informatics program that emphasizes both cognitive science as well as traditional medical informatics.

Acknowledgement

I would like to thank Vimla L. Patel and two anonymous reviewers for valuable comments and suggestions on an early draft. Correspondence and request for reprints should be sent to Jijie Zhang, School of Health Information Sciences, University of Texas Health Sciences Center at Houston, 7000 Fannin, Suite 600, Houston, TX, 77030, USA.

References

1. Coiera, E. (1997). Guide to medical informatics, the Internet and telemedicine. London: Arnold.
2. Huff, S. M., & Carter, J. S. (1999). A characterization of terminology models, clinical templates, message models, and other kinds of clinical information models. In C. Chute (Ed.), Proceedings of International Medical Informatics Association--Working Group 6 Meeting (Health Concept Representation and Natural Language Processing).
3. Shortliffe, E. H., Perreault, L. E., Wiederhold, G., & Fagan, L. M. (2001). Medi-

- cal informatics: Computer applications in health care and biomedicine. (2nd Edition). Springer: New York.
4. Van Bommel, J. H., & Musen, M. A. (1997). Handbook of medical informatics. Springer: New York.
 5. Cimino, J. J. (1997). Desiderata for controlled medical vocabularies in the twenty-first century. IMIA WG6 Conference, Jacksonville, Florida.
 6. Chute, C. G., Cohn, S. P., Campbell, K. E., Oliver, D. E., & Campbell, J. R. (1996) The content coverage of clinical classifications. *JAMIA*, 3 (3), 224-233.
 7. Patel, V. L., Kushniruk, A. W. (1997). Understanding, navigating and communicating knowledge: Issues and challenges. IMIA WG6 Conference, Jacksonville, Florida.
 8. Patel, V. L., Arocha, J. F., & Kushniruk, A. W. (2002). EMR: Re-engineering the organization of health information. *Journal of Biomedical Informatics*.
 9. Rumelhart, D. E. & Norman, D. A. (1988). Representation in memory. In R. C. Atkinson, R. J. Herrnstein, G. Lindzey, & R. D. Luce (Ed.), *Stevens' Handbook of Experimental Psychology*. New York: Wiley.
 10. Zhang, J. (1997). The nature of external representations in problem solving. *Cognitive Science*, 21, 179-217.
 11. Zhang, J. (2000). External representations in complex information processing tasks. In A. Kent (Ed.), *Encyclopedia of Library and Information Science* (Vol. 68, pp. 164-180). New York: Marcel Dekker, Inc.
 12. Zhang, J., & Norman, D. A. (1994). Representations in distributed cognitive tasks. *Cognitive Science*, 18, 87-122.
 13. Norman, D. A. (1988). *The psychology of everyday things*. New York: Basic Books.
 14. Hammond, W. E., & Cimino, J. J. (2001). Standards in medical informatics. In E. H. Shortliffe, L. E. Perreault, G. Wiederhold, and L. M. Fagan: *Medical informatics – Computer applications in health care and biomedicine*. Springer: New York.
 15. World Health Organization. (1977). Ninth Edition, International Classification of diseases index. Manual for the international statistical classification of diseases. Geneva: The World Health Organization.
 16. Zhang, J., & Norman, D. A. (1995). A representational analysis of numeration systems. *Cognition*, 57, 271-295.
 17. Barsalou, L. W., & Medin, D. L. (1989). Concepts: Fixed definitions or dynamic context-dependent representation? *Cahiers de Psychologie Cognitive*, 6, 187-202.
 18. Keil, F. C. (1986). The acquisition of natural kind and artifact terms. In W. Demopoulos & A. Marras (Eds.), *Language learning and concept acquisition*. Norwood, NJ: Ablex.
 19. Kintsch, W. (20021). The potential of Latent Semantic Analysis for machine grading of clinical case summaries. *Journal of Biomedical Informatics*.
 20. Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, 92, 289-316.
 21. Newell, A. (1991). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
 22. Rumelhard, D. E., Smolensky, P., McClelland, J. L., & Hinton, G. E. (1986). Schemata and sequential thought processes in PDP models. In J. L. McClelland & D. E. Rumelhard (Eds.), *Parallel and distributed processing: Explorations in the microstructure of cognition* (Vol. 2). Cambridge, MA: MIT Press.
 23. Schank, R. C., Collins, G. C., & Hunter, L. E. (1986). Transcending inductive

- category formation in learning. *Behavioral and Brain Sciences*, 9, 639-686.
24. Smith, E. E., & Medin, D. L. (1981). *Categories and concepts*. Cambridge, MA: Harvard University Press.
 25. Rosch, E. H., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, 7, 573-605.
 26. Medin, D. L., & Ross, B. H. (1992). *Cognitive Psychology*. New York : Harcourt Brace Jovanovich.
 27. Smith, E. E., Shoben, E. J., & Rips, L. J. (1974). Structure and process in semantic memory: A featural model form semantic decisions. *Psychological Review*, 81, 214-241.
 28. Medin, D. L., & Smith, E. E. (1984). Concepts and concept formation. In M. R. Rosenzweig & L. W. Porter (Eds.), *Annual Review of Psychology*, 35, 113-138.
 29. Hintzman, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review*, 93, 411-428.
 30. Medin, D. L., & Schaffer, M. M. (1978). A context theory of classification learning. *Psychological Review*, 85, 207-238.
 31. Rosch, E. H., Mervis, C. B., Johnson, D. M., & Bayes Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, 8, 382-439.
 32. McClelland, J. L., Rumelhart, D. E., & The PDP Research Group. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition. Volume 2: Psychological and biological models (Vol. 2)*. Cambridge, MA: MIT Press
 33. Rumelhart, D. E., McClelland, J. L., & The PDP Research Group. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition. Volume 1: Foundations (Vol. 1)*. Cambridge, MA: MIT Press
 34. Chi, M. T. H., Feltovich, P. J., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science*, 5, 121-152.
 35. Luria, A. R. (1976). *Cognitive development: its cultural and social foundations*. Cambridge, MA: Harvard University Press.
 36. Clancey, W. (1986). From GUIDON to NEOMYCIN and HERACLES in twenty short lessons. *AI Magazine*, 7, 40-60.
 37. Shortliffe, E. H. (1986). Medical expert systems: Knowledge tools for physicians. *Western Journal of Medicine*, 145, 830-839.