

Partitioning the UMLS Semantic Network

Zong Chen, Yehoshua Perl, Michael Halper, James Geller, Huanying Gu

Abstract—The Unified Medical Language System (UMLS) integrates many well established biomedical terminologies. The UMLS’s Semantic Network (SN) can help orient users to the vast knowledge content of the UMLS Metathesaurus via its abstract conceptual view. However, the SN itself is large and complex and may still be difficult to comprehend. Our technique partitions the SN into smaller meaningful units amenable to display on limited-sized computer screens. The basis for the partitioning is the distribution of the relationships within the SN. Three rules are applied to transform the original partition into a second, more cohesive, partition.

Keywords—UMLS, Semantic Network, IS-A Relationship, Partition, Semantic Type, Orientation, Views, Cohesive Partition

I. INTRODUCTION

The Unified Medical Language System (UMLS) [1], [2], [3], [4] combines many well established medical informatics terminologies into a unified knowledge representation system. The UMLS can be used to overcome problems caused by discrepancies in different terminologies [5], [6]. However, the UMLS’s enormous size and complexity (730,000 concepts in the Metathesaurus (META) [7]) can pose serious comprehension problems for potential users [8].

The UMLS’s Semantic Network (SN) helps to orient users [9] to the vast knowledge content of META. The SN is composed of a set of 134 semantic types which define high-level abstractions for sets of concepts from META [10], [11], [12]. Each concept in META is assigned to one or more semantic types in the SN. Overall, the semantic types are arranged in a hierarchy of IS-A relationships. In addition, there are 54 other kinds of (semantic) relationships which connect semantic types.

However, the SN abstract view of META can still be too large and difficult for comprehension. A convenient way for a user to get oriented to a large knowledge-base is by studying a diagram. For such a diagram to be easily comprehensible, it should fit on a computer screen and thus contain a limited number of nodes along with their interconnections. Fig. 1, showing only a quarter of the SN, is already difficult to comprehend, and it displays neither the incoming relationships nor the inherited relationships of the semantic types.

In this paper, we concentrate on providing comprehensible access to the SN through simpler views which fit easily onto a single screen. Such a need is even more urgent in light of a refined object-oriented database representation of SN, containing 1,296 classes which we created as an extension to the SN [13]. Specifically, we present a technique for partitioning the SN based on its relationship configuration. The outcome of our technique is a partition of the SN into sets of semantic types, called *semantic-type groups*. Considering some enhancements leads to a revised methodology that partitions the SN into *cohesive collections of semantic types*.

We note the existence of efficient algorithms for partitioning a tree structure according to various criteria such as Max-Min [14], Min-Max [15], *etc.* However, such quantitative criteria do not fit our purpose of obtaining cohesive units of semantically

related semantic types, each fitting a subject area. Although we apply some structural measures in our partitioning, they are related to semantics and result in a partition with the desired cohesiveness (as shown in [16]). Due to the need for semantics, zoom-like partitioning methods are not successful either.

A study was conducted in [16] to measure the meaningfulness of such a partition by comparing it to experts’ partitions of the SN, done according to semantic considerations. In [16], the collections of the cohesive partition also served to define various partial views of SN, resulting in a powerful viewing mechanism for the SN.

Section 2 proposes a method to partition the SN based on the structure of the relationships of its semantic types. Section 3 defines three rules to enhance the structural partition resulting in the *cohesive partition*. Section 4 contains the conclusions.

II. STRUCTURAL PARTITIONING

In the SN, the IS-A hierarchy supports the inheritance of the semantic relationships among semantic types. When two semantic types are linked via IS-A, the child semantic type inherits all the relationships defined for the parent semantic type. For example, the semantic type **Activity** IS-A **Event** and therefore inherits **Event**’s *issue_in* relationship.¹

The UMLS provides two additional modeling features that affect the inheritance of relationships. The first allows a newly introduced relationship to be designated as “defined but not inherited” (“DNI”), which means that the relationship is not inherited by any of the children of the semantic type that is introducing it. The second feature, called “blocking,” nullifies the definition of an inherited relationship.

The SN’s IS-A hierarchy has two roots, **Entity** and **Event**. We will demonstrate our technique on that portion of the SN rooted at **Event** (Fig. 1), with 35 semantic types, 34 IS-A relationships, and 134 (semantic) relationships. (Note that in Fig. 1, rectangles represent semantic types, IS-A relationships are represented by bold arrows, and other relationships appear as labeled thin arrows. A semantic type appearing outside the scope of the figure is denoted as a circle labeled with “?” inside. Also, the names of some relationships are written as numbers and listed in the legend of the figure.)

For a partitioning of the semantic types into groups to be effective for comprehension purposes, each group should have a unifying theme. That is, each group should be a logical unit composed of similar semantic types. Comprehension of such uniform sets is easier than comprehension of non-uniform sets.

Our partitioning technique is based on the distribution of the relationships among the semantic types of the SN. From now on, whenever we use “relationship” we mean a semantic relationship rather than IS-A. A relationship is introduced at a given semantic type and inherited by all its descendants (unless the inheritance is interrupted by the DNI designation or blocking). E.g., all descendants of **Phenomenon** or **Process** inherit *result_of*, which is introduced at that point. In Fig. 1, when a semantic type inherits a relationship from its parent but the target semantic type is refined, we show the inherited rela-

Zong Chen, Yehoshua Perl, and James Geller are at the CS Department, New Jersey Institute of Technology, Newark, NJ 07102. Michael Halper is at the Department of Mathematics & Computer Science, Kean University, Union, NJ 07083. Huanying Gu is at the Department of Health Informatics, UMDNJ, Newark, NJ 07103. Corresponding Author: James Geller: geller@homer.njit.edu

¹Let us note some typographical conventions used throughout the paper: A semantic type will be written in a bold font. The name of a semantic relationship will be written in italicized lowercase letters.

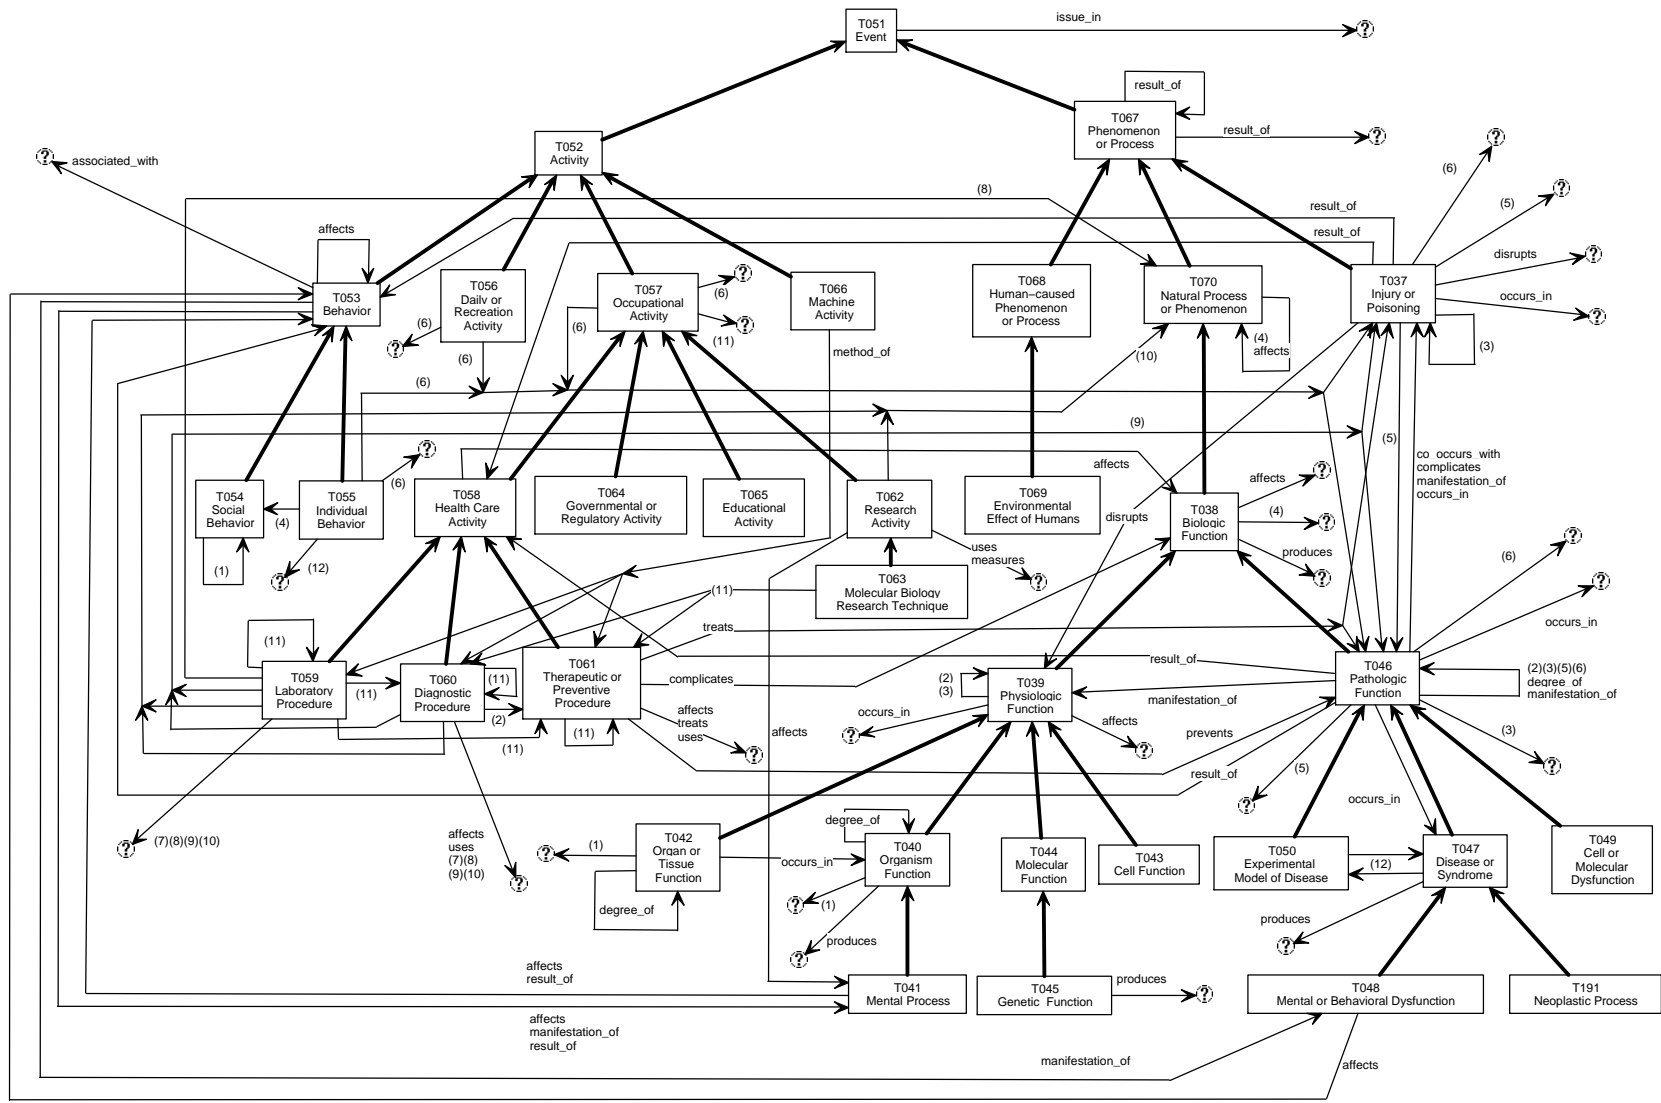


Fig. 1. Event portion of the UMLS SN

- (1) conceptual_part_of (DNI)
- (2) precedes
- (3) co occurs with
- (4) process of
- (5) complicates
- (6) associated_with
- (7) analyzes
- (8) assesses_effect_of
- (9) diagnoses
- (10) measures
- (11) method_of
- (12) conceptually_related_to (DNI)

tionship explicitly. For example, **Organ or Tissue Function** inherits the relationship *occurs_in*, defined at its parent **Physiologic Function**. However, it has a new target, **Organism Function**.

We focus on the relationships because of their overall definitional importance. In fact, we define the “structure” of a semantic type as the set of its defined relationships, whether they be introduced directly or inherited. Two semantic types are “structurally identical” if they both have the exact same set of relationships defined for them. The identical nature of their relationship structures suggests that they bear a close resemblance in meaning. It is therefore justified to group them together along that dimension of similarity to form unified logical units: All semantic types exhibiting the exact same set of relationships are grouped together. See [17], [18] for an example of using structural similarity to group concepts of the MED (Medical Entity Dictionary) terminology for the purpose of producing a terminology schema [19]. For another structural technique for partitioning a vocabulary, see [20].

Definition (Semantic-type group): A *semantic-type group* is an abstract conceptual entity comprising the set of all semantic types with the exact same set of relationships. ■

Definition (Root of a semantic-type group): A semantic type is a root of a semantic-type group if none of its parents belong to that semantic-type group. ■

Clearly, a semantic type which introduces a new relationship will be a root of its semantic-type group. Most, but not all, semantic-type groups have unique roots. If a semantic-type group has a unique root, then all other semantic types in the group are its descendants, and thus are more specialized semantic types of the root semantic type.

Taken altogether, the semantic-type groups of the SN form a partition: Every semantic type must be in one semantic-type group, and, in fact, that semantic-type group is unique.

The semantic type **Event**, the root of this portion of the SN hierarchy, introduces the relationship *issue_in* and therefore starts a new semantic-type group. **Activity** inherits **Event**’s *issue_in* relationship and does not introduce any new relationships of its own. Hence, it belongs to **Event**’s semantic-type group. In contrast, the other child **Phenomenon or Process** introduces a new relationship *result_of* and starts another semantic-type group. (See Fig. 2 where semantic-type groups with more than one member are enclosed in dashed bubbles.)

Overall, the event hierarchy of the SN is partitioned into 21 semantic-type groups, as shown in Fig. 2. For the entire SN, there are 71 semantic-type groups. Of these, 47 contain just one semantic type. (We call such groups “singletons.”) 11 groups have two semantic types; five groups have three semantic types; three groups have four semantic types; two groups have five semantic types; one group has six semantic types, and one other has eight. Finally, there is one group with fourteen.

Note that a semantic type s which introduces a DNI relationship is a root of a group since its parent semantic type’s structure does not contain this relationship. On the other hand, this relationship is not contained in the structure of any child semantic type of s because of the lack of inheritance due to the DNI designation. We call such a semantic type a *DNI-root*.

III. COHESIVE PARTITIONING

For an effective partitioning of the SN, the groups of semantic types have to be not just uniform in their structure but also cohesive. For a group of semantic types to be cohesive, it should have a unique root, i.e., one semantic type which all other semantic types in the group are descendants of. The co-

hesiveness is a result of the fact that each one of the semantic types in the group is a specialization of the unique root. Hence, by naming the semantic-type group after the root, this name properly reflects the overarching semantics of the group. As we see in Fig. 2, most of the semantic-type groups have unique roots. There are only a few that do not, and these are said to be non-cohesive. This phenomenon shows that uniform structure groups tend to be cohesive most of the time, but not all the time. Since cohesiveness is also important for comprehension, we will provide, in this section, rules to convert the structural partitioning into a cohesive partitioning. For this conversion, we will need to make some trade-offs, meaning some cohesive groups will lose their structural uniformity. However, they will still have approximate structural uniformity.

Another problem with the structural partitioning is the large number of singletons, which do not help comprehension. Thus, we will provide a rule to add singletons to other semantic-type groups to minimize the number of singletons in the partition. Again, this implies creating groups which are not structurally uniform, since those singletons were created due to structural differences. The rule that we will provide will, nevertheless, ensure that the new groups have approximate structural uniformity.

The cohesive partition which will emerge from applying our rules to the structural partition will be based on *semantic-type collections*. Each semantic-type collection is an abstract conceptual entity representing a set of semantic types in the SN. Each will have a unique root and will thus be cohesive. When a semantic-type collection is also a semantic-type group, it will be structurally uniform. Otherwise, it will have approximate structural uniformity.

Rule 1: Each semantic-type group with a non-leaf unique root becomes a semantic-type collection. ■

The second rule deals with “leaf” semantic types which are semantic types without children.

Rule 2: If a leaf semantic type is in a singleton in the structural partitioning, and its parent semantic-type group does not have a DNI-root, then it is added to the semantic-type collection which contains its parent. ■

Note, e.g., that the singletons containing the leaves **Social Behavior** and **Individual Behavior** (Fig. 2) are combined with the singleton containing **Behavior** to produce a new semantic-type collection with three members (Fig. 3). Applying Rule 2 helps to merge many singleton semantic-type groups into larger semantic-type collections. We still allow a non-leaf singleton in the partitioning since it may play a role as a branching point.

Let us define the “structure” of a semantic-type collection to be the set of relationships of its root. The structure of a leaf added to a semantic-type collection is equal to the union of the semantic-type collection’s structure with any relationships introduced by the leaf. (This is true in the case where none of the relationships of the leaf’s parent semantic-type collection are declared DNI.) It will be noted that no other semantic-type collection can have the structure of that leaf. Furthermore, several leaves added to a semantic type collection will have all the relationships of the root of the semantic type collection in common. Thus, although the semantic types in a semantic-type collection do not always have uniform structure, that structure is approximately uniform and unique.

The children, **Social Behavior** and **Individual Behavior**, of **Behavior** were originally in singletons. On one hand, they inherit all relationships of **Behavior**. On the other hand, they are structurally different from their parent. **Social Behavior**

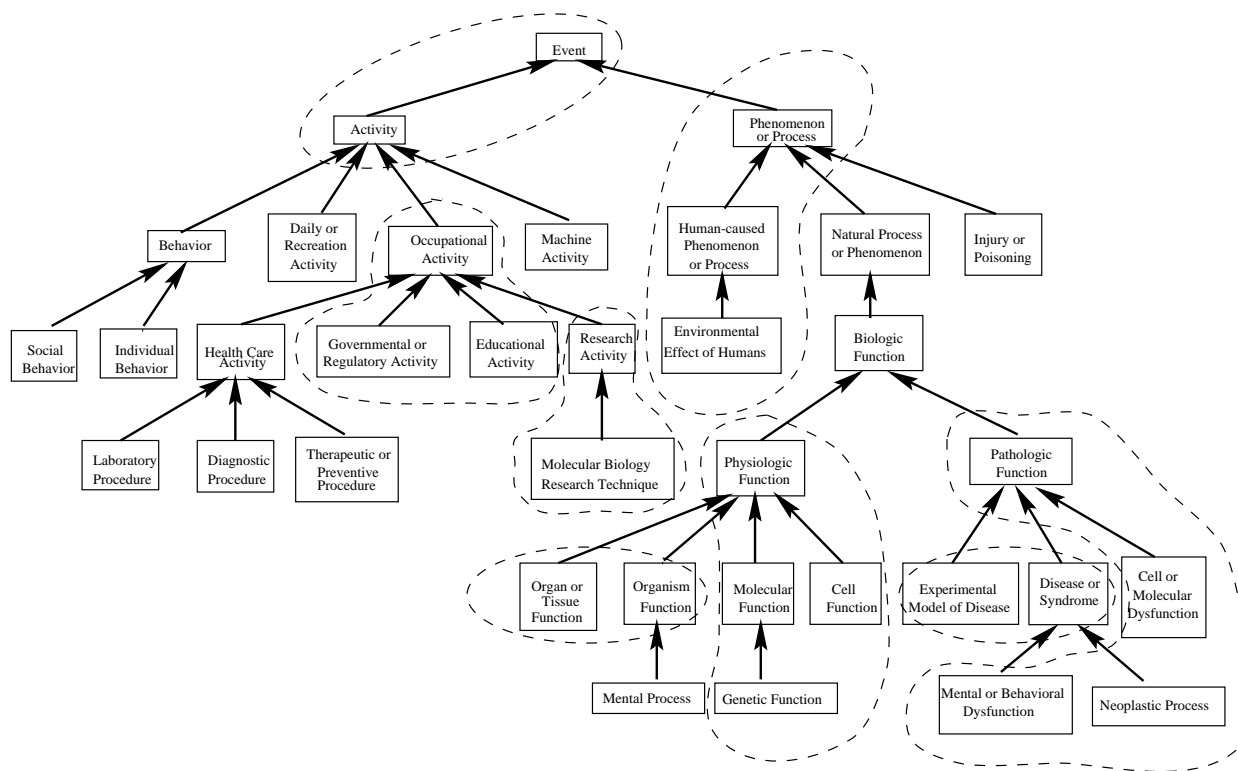


Fig. 2. Structural partition consisting of semantic-type groups

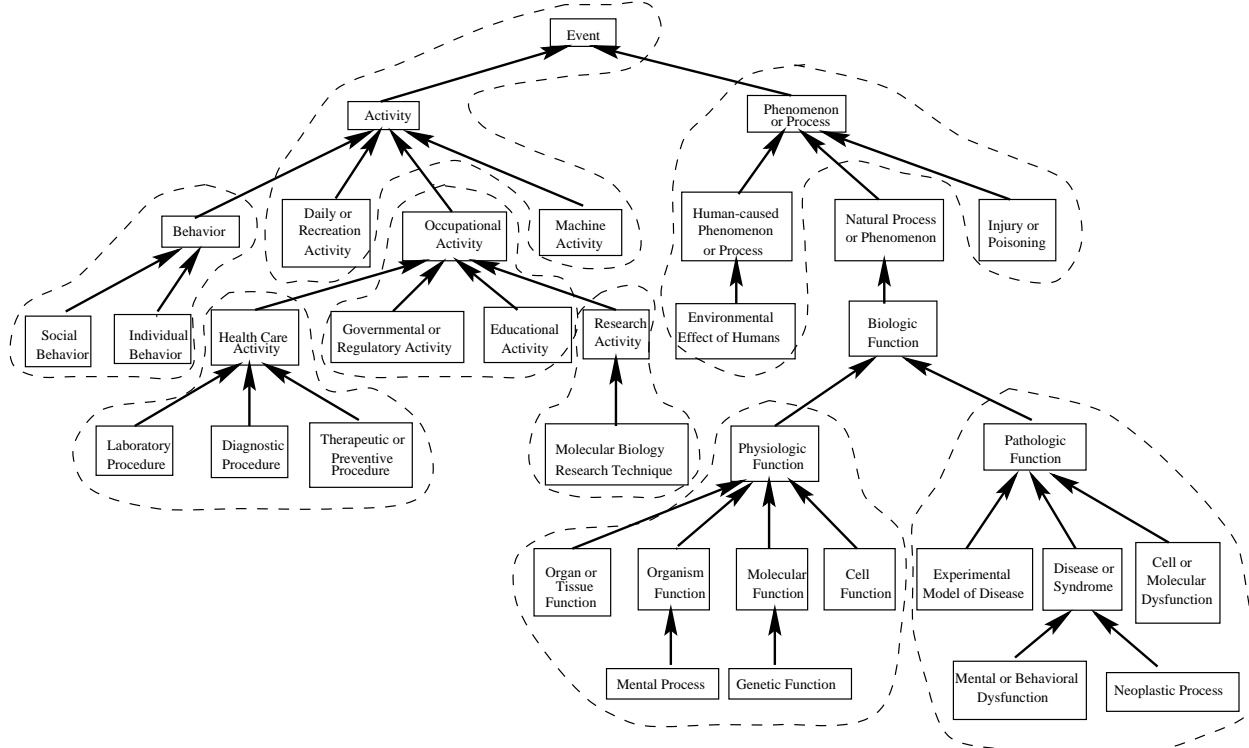


Fig. 3. Cohesive partition consisting of semantic-type collections

introduces a new relationship *conceptual_part_of* directed at itself. **Individual Behavior** introduces *process_of* directed at **Social Behavior**, and two other relationships. However, their structure is more similar to that of **Behavior** than it is to that of other semantic-type collections.

We are now turning our attention to cases of semantic-type groups with multiple roots. In our discussion, we will concen-

trate on those appearing in Fig. 2. One example contains the sibling semantic types **Organ or Tissue Function** and **Organism Function**. Both inherit all relationships of their parent **Physiologic Function** and introduce the new relationship *degree_of*.

Another example centers around **Experimental Model of Disease** and **Disease or Syndrome**, each of which defines the

new relationship *conceptually_related_to* directed at the other. Both semantic types designate this relationship “DNI,” meaning that it is not inherited by any of their respective children. Since the children of **Disease or Syndrome** do not introduce any new relationships, the subtree rooted at **Pathologic Function** is partitioned into two semantic-type groups. One includes **Experimental Model of Disease** and **Disease or Syndrome**. The other includes **Pathologic Function**, **Cell or Molecular Dysfunction**, **Neoplastic Process**, and **Mental or Behavioral Dysfunction** (again, see Fig. 2). For such cases, we need to introduce an extra rule.

Rule 3: Let the semantic types A_1, A_2, \dots, A_n ($n \geq 2$) be roots of the same semantic-type group G of the structural partitioning. Add all semantic types of G to the semantic-type collection of their lowest common ancestor in the IS-A hierarchy, assuming the root of that semantic-type collection is not a DNI-root. ■

For example, **Organ or Tissue Function** and **Organism Function** join the semantic-type collection rooted at **Physiologic Function** (Fig. 3) based on Rule 3. Then, by Rule 2, the semantic type **Mental Process** also joins the same semantic-type collection. Hence, all the descendants of **Physiologic Function** belong now to its semantic-type collection.

Fig. 3 shows the SN’s event hierarchy after the cohesive partitioning technique has been applied to identify the final semantic-type collections. The semantic-type collections *Natural Process or Phenomenon* and *Biologic Function* are still singletons. *Biologic Function* is the branching point into the *Physiologic Function* collection and the *Pathologic Function* collection. There are 10 semantic-type collections in the figure. In the entire SN, there are 28 semantic-type collections. In Table 1, we list the 28 semantic-type collections of the SN alphabetically, with the number and list of semantic types in each collection. The average size of a semantic-type collection is 4.786. In Table 2, we list the distribution of the sizes of the 28 semantic-type collections.

Based on this partition, we developed in [21] a metaschema which is a compact abstract level for the Semantic Network. Each node of the metaschema corresponds to a collection of the cohesive partition. An alternative partition of SN into 14 groups is presented by McCray *et al.* in [22]. Each group of [22] represents a subject area. Their partition is designed to conform to six criteria: validity, parsimony, completeness, exclusivity, naturalness, and utility. However, some of their groups are disconnected. This contradicts the property of validity, as noted by the authors themselves. Due to this fact, the partition of [22] does not lend itself to the extraction of a natural metaschema. This stands in contrast to the partition that we presented. The metaschema, together with the partial views of [16], help the user to obtain an orientation in the collection of semantic types and an understanding of the interactions among them.

IV. CONCLUSIONS

We have presented a partitioning method, specifically designed for the Semantic Network of the UMLS, that relies on structural similarity between semantic types, combined with a step to eliminate singleton leaves and a step to find unique roots for each partition. The resulting partition contains 28 semantic type collections of sizes that are easily displayable on one screen. In [16], we present an evaluation study that compares our cohesive partition with the results that human experts obtain when partitioning the (same part of the) UMLS Semantic Network.

Acknowledgment

We thank Dr. James Cimino for commenting on an earlier draft of this paper.

REFERENCES

- [1] K. E. Campbell, D. E. Oliver, and E. H. Shortliffe. The Unified Medical Language System: Toward a collaborative approach for solving terminologic problems. *JAMIA*, 5(1):12–16, 1998.
- [2] B. L. Humphreys and D. A. B. Lindberg. Building the Unified Medical Language System. In *Proc. Thirteenth Annual Symposium on Computer Applications in Medical Care*, pages 475–480, Washington, DC, November 1989.
- [3] B. L. Humphreys, D. A. B. Lindberg, H. M. Schoolman, and G. O. Barnett. The Unified Medical Language System: An informatics research collaboration. *JAMIA*, 5(1):1–11, 1998.
- [4] D. A. B. Lindberg, B. L. Humphreys, and A. T. McCray. The Unified Medical Language System. *Methods of Information in Medicine*, 32:281–291, 1993.
- [5] J. J. Cimino. Review paper: Coding systems in health care. *Methods of Information in Medicine*, 35:273–284, 1996.
- [6] B. L. Humphreys and D. A. B. Lindberg. The Unified Medical Language System project: A distributed experiment in improving access to biomedical information. *Methods of Information in Medicine*, 7(2):1496–1500, 1992.
- [7] US Dept. of Health and Human Services, NIH, National Library of Medicine. *Unified Medical Language System (UMLS)*, 2000.
- [8] M. S. Tuttle, D. D. Sherertz, M. S. Erlbaum, et al. Adding your terms and relationships to the UMLS metathesaurus. In P. D. Clayton, editor, *Proceedings of the Fifteenth Annual SCAMC*, pages 219–223, Washington, D.C., 1991.
- [9] A. T. McCray and S. J. Nelson. The representation of meaning in the UMLS. *Methods of Information in Medicine*, 34:193–201, 1995.
- [10] A. T. McCray. UMLS semantic network. In *Proceedings of the Thirteenth Annual SCAMC*, pages 503–507, 1989.
- [11] A. T. McCray. Representing biomedical knowledge in the UMLS semantic network. In N. C. Broering, editor, *High-Performance Medical Libraries: Advances in Information Management for the Virtual Era*, pages 45–55. Meckler, Westport, CT, 1993.
- [12] A. T. McCray and W. T. Hole. The scope and structure of the first version of the UMLS semantic network. In *Proceedings of the Fourteenth Annual SCAMC*, pages 126–130, 1990.
- [13] H. Gu, Y. Perl, J. Geller, M. Halper, L. Liu, and J. J. Cimino. Representing the UMLS as an OODB: Modeling issues and advantages. *JAMIA*, 7(1):66–80, Jan/Feb 2000.
- [14] Y. Perl and S.R. Schach. Max-Min tree partitioning. *JACM*, 28:5–15, 1981.
- [15] R.I. Becker, Y. Perl, and S.R. Schach. A shifting algorithm for Min-Max tree partitioning. *JACM*, 29:58–76, 1982.
- [16] J. Geller, Y. Perl, M. Halper, and Z. Chen. Evaluation and application of a semantic network partition. *IEEE Transactions on Information Technology in Biomedicine*, in this issue.
- [17] H. Gu, M. Halper, J. Geller, and Y. Perl. Benefits of an OODB representation for controlled medical terminologies. *JAMIA*, 6(4):283–303, July/Aug 1999.
- [18] L. Liu, M. Halper, J. Geller, and Y. Perl. Controlled vocabularies in OODBs: Modeling issues and implementation. *Distributed and Parallel Databases*, 7(1):37–65, January 1999.
- [19] J. J. Cimino, P. D. Clayton, G. Hripcsak, and S. Johnson. Knowledge-based approaches to the maintenance of a large controlled medical terminology. *JAMIA*, 1(1):35–50, 1994.
- [20] H. Gu, J. Geller, L. Liu, and M. Halper. Using a similarity measurement to partition a vocabulary of medical concepts. In *Proc. 10th Int’l Conference on Database and Expert Systems Applications: DEXA’99*, Florence, Italy, August 1999. Published as: *Lecture Notes in Computer Science*, 1677, Springer-Verlag, New York, pages 712–723.
- [21] M. Halper, Z. Chen, J. Geller, and Y. Perl. A metaschema of the UMLS based on a partition of its semantic network. In *Proceedings of the 2001 AMIA Annual Symposium*. In press.
- [22] A. T. McCray, A. Burgun, and O. Bodenreider. Aggregating UMLS semantic types for reducing conceptual complexity. To appear in: *Proceedings of Medinfo 2001*.

TABLE I
SEMANTIC-TYPE COLLECTION LIST

<i>Collection</i>	<i>Sz.</i>	<i>Semantic Types in Collection</i>
Anatomical Abnormality	3	Anatomical Abnormality; Congenital Abnormality; Acquired Abnormality
Anatomical Structure	2	Anatomical Structure; Embryonic Structure
Animal	9	Animal; Invertebrate; Vertebrate; Amphibian; Bird; Fish; Reptile; Mammal; Human
Behavior	3	Behavior; Social Behavior; Individual Behavior
Biologic Function	1	Biologic Function
Biologically Active Substance	7	Biologically Active Substance; Receptor; Vitamin; Enzyme; Neuroreactive Substance or Biogenic Amine; Hormone; Immunologic Factor
Chemical	16	Chemical; Chemical Viewed Functionally; Hazardous or Poisonous Substance; Inorganic Chemical; Biomedical or Dental Material; Element, Ion or Isotope; Indicator, Reagent or Diagnostic Aid; Carbohydrate; Chemical Viewed Structurally; Organic Chemical; Organophosphorus Compound; Steroid; Eicosanoid; Amino Acid, Reptide, or Protein; Lipid; Nucleic Acid, Nucleoside, or Nucleotide
Entity	8	Entity; Physical Object; Concept Entity; Group Attribute; Language; Intellectual Product; Classification; Regulation or Law
Event	4	Event; Activity; Daily or Recreation Activity; Machine Activity
Finding	3	Finding; Lab or Test Result; Sign or Symptom
Fully Formed Anatomical Structure	6	Fully Formed Anatomical Structure; Cell; Cell Component; Tissue; Gene or Genome Body Part, Organ, or Organ Component;
Group	6	Group; Professional or Occupational Group; Population Group; Family Group; Age Group; Patient or Disabled Group
Health Care Activity	4	Health Care Activity; Laboratory Procedure; Dianostic Procedure; Therapeutic or Preventive Procedure
Idea or Concept	14	Idea or Concept; Functional Concept; Body System; Temporal Concept; Qualitative Concept; Quantitative Concept; Spatial Concept; Geographic Area; Body Location or Region; Molecular Sequence; Carbohydrate Sequence; Amino Acid Sequence; Body Space or Junction; Nucleotide Sequence
Manufactured Object	4	Manufactured Object; Medical Device; Research Device; Clinical Drug
Natural Phenomenon or Process	1	Natural Phenomenon or Process
Occupation or Discipline	2	Occupation or Discipline; Biomedical Occupation or Discipline
Occupational Activity	3	Occupational Activity; Educational Activity Governmental or Regulatory Activity;

Zong Chen received his B.S. degree in Electrical Engineering at the East China Institute of Technology, Nanjing, China (1992), and his M.S. degree in Electrical Engineering at the Southeast University, Nanjing, China (1997). He is currently working on his Ph.D. degree in Computer Science at NJIT. He has published several papers in journals and conferences on topics in medical vocabularies.

Yehoshua Perl received his Ph.D. degree in CS in 1975 from the Weizmann Institute of Science, Israel. Since 1985, he has been in the Computer and Information Science Department at NJIT, where he was appointed professor in 1987. Dr. Perl received the Harlan Perlis Research Award of NJIT in 1996. Dr. Perl is the author of more than 100 papers in international journals and conferences. His publications are in object-oriented databases, design and analysis of algorithms, data structures, design of networks, sorting networks, graph theory, and data compression. His current research deals with modeling medical vocabularies using object-oriented databases.


larities using object-oriented databases.

TABLE I (CONTINUED)
SEMANTIC-TYPE COLLECTION LIST

Organism	6	Organism; Archaeon; Virus; Bacterium; Fungus; Rickettsia or Chlamydia
Organism Attribute	2	Organism Attribute; Clinical Attribute
Organization	4	Organization; Health Care Related Organization; Professional Society; Self-help or Relief Organization
Pathologic Function	6	Pathologic Function; Experimental Model of Disease; Cell or Molecular Dysfunction; Disease or Syndrome; Mental or Behavioral Dysfunction
Pharmacologic Substance	2	Pharmacologic Substance; Antibiotic
Phenomenon or Process	4	Phenomenon or Process; Injury or Poisoning; Human-caused Phenomenon or Process; Environmental Effect of Humans
Physiologic Function	7	Physiologic Function; Organ or Tissue Function; Organism Function; Mental Process; Molecular Function; Genetic Function; Cell Function
Plant	2	Plant; Alga
Research Activity	2	Research Activity; Molecular Biology Research Technique
Substance	3	Substance; Body Substance; Food


TABLE II
DISTRIBUTION OF SIZES OF SEMANTIC-TYPE COLLECTIONS

Size of the collection:	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Number of groups of that size:	2	6	5	5	0	4	2	1	1	0	0	0	0	1	0	1



Michael Halper received the CS B.S. degree (with honors) from NJIT (1985), the CS M.S. degree from Fairleigh Dickinson University (1987) and the CS Ph.D. degree from NJIT (1993). He was the recipient of a Garden State Graduate Fellowship from the State of New Jersey. Dr. Halper is an associate professor of computer science at Kean University. His research interests include conceptual and object-oriented modeling, part-whole modeling, extensible data models, object-oriented database systems, and medical informatics. Dr. Halper has published numerous papers in international journals and conferences. He is a mem-

ber of The Honor Society of Phi Kappa Phi.



Huanying Gu received her B.S. degree in Computer Science from Huazhong University of Science and Technology in 1984 and her M.S. degree in Computer Science from Nanjing University of Science and Technology in 1989, both in China. She received a Ph.D in Computer Science from New Jersey Institute of Technology in 1999. Huanying Gu has been with Health Informatics Department in University of Medicine and Dentistry of New Jersey as an assistant professor since 1999. Her research interests include object-oriented modeling, object-oriented databases, knowledge representation, controlled medical termi-

nologies, and biomedical informatics.



James Geller received an Electrical Engineering Diploma from the Technical University Vienna, Austria (1979), an M.S. degree (1984) and a Ph.D. degree (1988) in Computer Science from SUNY at Buffalo. James Geller is professor in the Computer Science Department at NJIT where he is also Director of the OODB & AI Laboratory and Vice Chair of the M.S. and Ph.D. programs in Biomedical Informatics. Dr. Geller has published numerous journal and conference papers in knowledge representation, parallel artificial intelligence, medical informatics, and object-oriented databases. His current research interests concentrate

on object-oriented modeling of medical vocabularies, and on Web mining.