MEDINFO 2017: Precision Healthcare through Informatics A.V. Gundlapalli et al. (Eds.) © 2017 International Medical Informatics Association (IMIA) and IOS Press. This article is published online with Open Access by IOS Press and distributed under the terms of the Creative Commons Attribution Non-Commercial License 4.0 (CC BY-NC 4.0). doi:10.3233/978-1-61499-830-3-978

Summarizing an Ontology: A "Big Knowledge" Coverage Approach

Ling Zheng, Yehoshua Perl, Gai Elhanan, Christopher Ochs, James Geller, Michael Halper

College of Computing, New Jersey Institute of Technology, Newark, NJ 07102-1982, USA

Abstract

Maintenance and use of a large ontology, consisting of thousands of knowledge assertions, are hampered by its scope and complexity. It is important to provide tools for summarization of ontology content in order to facilitate user "big picture" comprehension. We present a parameterized methodology for the semi-automatic summarization of major topics in an ontology, based on a compact summary of the ontology, called an "aggregate partial-area taxonomy", followed by manual enhancement. An experiment is presented to test the effectiveness of such summarization measured by coverage of a given list of major topics of the corresponding application domain. SNOMED CT's Specimen hierarchy is the test-bed. A domain-expert provided a list of topics that serves as a gold standard. The enhanced results show that the aggregate taxonomy covers most of the domain's main topics.

Keywords:

Big Knowledge; Topic Coverage; Ontology Summarization

Introduction

The Big Data to Knowledge (BD2K) initiative is expected to produce many knowledge items that can be expressed as assertions or as rules. However, orientation into large knowledge bases is a challenge by itself, the "Big Knowledge" challenge. Without some high-level mental representation about the kinds of content in a large knowledge base, effective use of the knowledge may be limited [1]. When an ontology surpasses thousands of assertions, even its curators are confronted with the problem of seeing the "big picture" of its content. The work of curators, in charge of developing and maintaining the ontology, relies on its comprehension. A topic of an ontology represented by a concept c is considered a major topic if c has a large number of descendants. The effectiveness of tools for summarizing and supporting users' comprehension of a large ontology can be measured by "Big Knowledge" coverage of a given list of major topics related to the corresponding domain. "Big Knowledge" coverage is defined as the percentage of the number of major topics in an ontology out of a given list, rather than by the percentage of the number of concepts out of a given list of concepts [2].

We summarize the "big picture" of an ontology by automatically deriving concept groups that represent major topics in a specific domain. The size of a topic, represented by a concept c in an ontology, is the number of c's descendants. The knowledge representation importance of topic c (which is different from its clinical importance) can be approximated by its size. The derivation of the topic-defining concepts of an ontology is based on its *aggregate partial-area taxonomy* ("*aggregate taxonomy*," for short) [3]. The aggregate taxonomy is a significantly smaller network that summarizes an ontology's concepts, providing a simplified view of the ontology. Aggregate taxonomies are based on *partial-area taxonomies* (*"taxonomies"* for short) [4], a compact summarization abstraction network developed, for example SNOMED CT [5].

In this paper, we demonstrate why *aggregate taxonomies* are better at summarizing an ontology than *taxonomies*, when measuring the effectiveness of the former by comparing with a gold standard list of major topics selected by a domain expert. We used SNOMED CT's *Specimen* hierarchy as a test-bed.

Background

Biomedical ontologies provide terminological support for Electronic Health Records (EHRs) [6], decision-support systems, natural language processing, data integration [7], etc. SNOMED CT is a comprehensive ontology organized into 19 hierarchies, such as *Clinical finding*, *Procedure* and *Specimen*. Its 317,057 active concepts (July 2015 release) are linked by *IS-A* relationships and more than 1.5 million attribute relationships (*'relationships*," for short). There are 1,620 concepts in the *Specimen* hierarchy.

We have developed *partial-area taxonomies* [4] to provide a compact view of the seven SNOMED CT hierarchies, including *Specimen*, that have outgoing relationships. We illustrate the process of the partial-area taxonomy derivation for an excerpt of the *Specimen* hierarchy (Figure 1(a)).

We start with the definition of the *area taxonomy*. We define an *area* as a set containing all concepts having the same set of outgoing relationships, without considering the ranges of those relationships. Areas are disjoint (i.e., each concept appears in only one area). Areas are named by their sets of relationships. Areas are represented by nodes of the area taxonomy. Figure 1(b) shows the four areas derived from the concepts in Figure 1(a). Areas are color-coded according to the cardinalities of their sets of relationships (i.e., all areas with the same number of relationships have the same color). For example, both areas {*Specimen source topography*} and {*Specimen substance*} in Figure 1(b) have one relationship, so both are colored green.

We define a *root* of an area as a concept having no parents in the area. Areas may have multiple roots. Area nodes are connected hierarchically, using *child-of* links, to form the area taxonomy. Those links are based on the configurations of area roots in the underlying ontology. More specifically, an area A is *child-of* an area B if a root in A has a parent in B. Figure 1(b) has four *child-of*'s derived in this manner.

A *partial-area taxonomy* ("taxonomy" for short) is derived from the area taxonomy as follows. We group a root and all its intra-area descendants into a *partial-area*, which is drawn as an embedded node of the area (see Figure 1(c)). If an area has



Figure 1-(a) Excerpt of 15 concepts from the Specimen hierarchy. Each set of concepts with the same set of relationships, called an area, is enclosed in a colored border, e.g., Body substance sample, Fluid sample, and Water specimen have one outgoing relationship Specimen substance. (b) Area taxonomy for (a). An area node is represented by a colored box labeled with its set of relationships and number of concepts. The three concepts having the Specimen substance relationship are now represented by an area named {Specimen substance}. Child-of links between areas are bold arrows. (c) Partial-area taxonomy for (a), composed of seven partial-areas, represented by embedded white boxes in their corresponding areas. Each is labeled with its root and number of concepts



Figure 2–An excerpt of the taxonomy for the Specimen hierarchy. Partial-areas are sorted (left-right and top-down) according to their size. The yellow partial-areas are the descendant partial-areas of the pink partial-area Specimen from trunk. That is, there is a path of child-of relationships from any yellow partial-area to Specimen from trunk

multiple roots, then its concepts are divided into multiple partial-areas to reflect their varied semantics. We label a partial-area node by its root. Partial-areas are not necessarily disjoint due to concepts with multiple parents/ancestors. Partial-areas are also connected by hierarchical *child-of* links to form the complete taxonomy. Figure 1(c) is the taxonomy for Figure 1(a). Figure 2 shows an excerpt of the taxonomy (with *child-of* omitted) for the entire *Specimen* hierarchy.

Partial-area taxonomies provide a structure-based summary of an ontology. Other ontology summarization techniques have been investigated outside of the biomedical domain, e.g., based on key concept identification by Li et al. [8].

Methods

First, we assume that concepts belonging to a given topic are all hierarchically related (i.e., they share a common ancestor concept c that represents and names the topic). That is, all the descendant concepts of a topic c belong to that topic since they are specification of c. And second, we assume that if

there are relatively more concepts for a topic then it is "more important." For example, there are 262 concepts related to digestive system specimens, but only 12 related to bone marrow specimens. We thus consider the topic "digestive system samples" as more important in SNOMED CT. We do not claim necessarily that it is more important clinically since this depends on clinical context.

Our approach for evaluating the automatically identified major topics is based on a gold standard list. A domain expert (GE) was asked to select a list of major topics for the specimen domain. (GE) is an MD with long experience in ontologies. A gold standard may also be derived from a published ontology of an authoritative organization. We did not find another ontology for specimens (e.g., in UMLS. For the sake of normalization and to simplify the eventual matching task), each chosen topic was semi-automatically mapped to a SNOMED CT concept in the *Specimen* hierarchy, utilizing UMLS synonyms. For example, the topic "Bone specimen" was mapped to the concept *Specimen from bone*.

One straightforward heuristic for identifying major topics in an typically general and cover high-level topics. For example, Specimen has 59 children (e.g., Biopsy sample and Blood specimen). However, among the 59 children, many would not be considered major topics (based on our second assumption above), since they have few descendants. For example, 13 of Specimen's children do not themselves have children (e.g., Muscle specimen). Nine have few children and no grandchildren (e.g., Fibroblast specimen has one child). Of the remaining 37 children, only 13 were in the major topic list of our domain expert, while another eight on that list were not children of Specimen (e.g., Stool specimen is a grandchild of Specimen). Hence, a better methodology for identifying major topics is required. In this study we address two questions: (1) how well do the partial-areas in a taxonomy match the topics in the list by a domain expert?, and (2) can we modify a taxonomy to automatically capture major topics in the ontology?

Partial-Area Taxonomies for Topic Identification

As the originators of parital-area taxonomies, we note that they were not designed for major topic identification, but for structure and content summarization. Indeed, the roots of partial-areas are not necessarily intuitive topics. The root is distinguished by the introduction of a new relationship type into the ontology, which may, but is not guranteed to, correlate with a major topic. Moreover, a partial-area may be small, and thus, may not define a broad topic. A taxonomy typically has many small partial-areas [3]. As a result, the taxonomy for a large ontology, although smaller by an order of magnitude than the ontology, can still fail to identify major topics. Metaphorically, the "forest" summary of the topics is not seen for the many small "trees" (see Figure 2).

Hence, a better solution for identifying major topics is to pick only the large partial-areas (with, e.g., dozens or more concepts). To illustrate these points, Figure 2 shows an excerpt of the Specimen taxonomy. Some concepts appear as (labels of) relatively large partial-areas. For example, Specimen from trunk (132), Specimen from head and neck structure (53), and Specimen from digestive system (50) from the area {Specimen source topography} are partial-areas with 50 or more concepts. However, all the seven large partial-areas account for only 536 Specimen concepts (33.1%). One may wonder about the topics of the other 66.9% of concepts.

Moving to medium-sized partial-areas with 20–49 concepts, we find eight partial-areas covering 218 (13.5%) concepts (e.g., *Blood specimen* (28) and *Soft tissue biopsy sample* (23)). Together, the large and medium partial-areas cover only 754 specimen concepts (46.5%). There are other problems with the summarization provided by the large/medium partial-areas. For example, all descendant partial-areas (yellow) of *Specimen from trunk* (pink) in Figure 2 contain refinements of this topic. They are in a separate partial-area because they have an extra relationship and appear in another area. For example, *Swab from abdomen* (13) has an additional *Specimen procedure* relationship. Overall, there are 201 partial-area descendants of *Specimen from trunk*, covering 551 concepts.

If we only focus on large and medium partial-areas, we are ignoring useful knowledge that is distributed among the many small partial-areas. Frequently, a large partial-area has many descendant small partial-areas. The concepts in these descendant partial-areas cover the same topic as the large parent/ancestor partial-area, but in more detail. Hence, they could also be summarized by the parent/ancestor partial-area.

Weighted Aggregate Partial-Area Taxonomies

We introduce an aggregation process that allows small partial-

areas to contribute to the identification of major topics. Large partial-areas are used as candidates for topics, while small ones are folded into their larger ancestor partial-area(s). Thus, the lost knowledge in small partial-areas is accounted for.

To address the inclusion of small partial-areas formally, a variation of a partial-area taxonomy based on an *adjustable minimum grouping threshold value* was introduced [3]. It is called the *aggregate partial-area taxonomy* (or just "aggregate taxonomy"). The derivation of an *aggregate taxonomy* begins with the selection of a threshold *b*, indicating the minimum size of a partial-area that will be included in the aggregate taxonomy. Then, using a *topological sort*, the aggregate taxonomy is generated by aggregating any partial-areas with sizes below *b* into their parent/ancestor partial-area(s) with sizes $\geq b$. (The root partial-area may be any size.) The nodes of this taxonomy are called *aggregate partial-areas*.

However, there is still another problem due to the structure of a *partial-area taxonomy*. We discovered that some major topics did not appear in the aggregate taxonomies at all, due to the small sizes of their partial-areas, in spite of having many small descendant partial-areas belonging to the same topic. For example, the partial-area capturing the topic *Endocrine sample* has only 10 concepts, but many more descendant concepts belong to this topic.

To overcome this difficulty, we defined an **aggregated weight** for each partial-area. This aggregated weight equals the sum of the size x of the partial-area itself and the sizes of all its descendant partial-areas smaller than x. In this way, the decision of which "small partial-areas" to eliminate from the aggregate taxonomies can now be based on the aggregated weight of a partial-area, rather than its size.

This new taxonomy is called *weighted aggregate partial-area taxonomy*. For example, the partial-area *Endocrine sample* (10) in Figure 3(a) does not appear in the aggregate taxonomy when b>10. However, its aggregated weight is 26, because it has 9 descendant partial-areas with fewer than 10 concepts, summarizing 16 descendant concepts. Therefore, the partial-area *Endocrine sample* will appear in the weighted aggregate partial-area taxonomy as long as $b \le 26$ (Figure 3(b)).



Figure 3– (a) An excerpt of 10 partial-areas. (b) Weighted aggregate partial-area for (a) with b=7 shown as a rounded white rectangle with its number of concepts in (), including all concepts from aggregated partial-areas, and the number of aggregated partial-areas in []

| - 981 |
|-------|
|-------|

| Topic | Concept | Partial-area | Weight | 1 | 5 | 10 | 15 | 20 | 25 | 30 |
|-------------------------------|--------------------------------|--|--------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| Blood specimen | Blood specimen | Blood specimen (28) | 43 | ~ | ~ | √ | ~ | ~ | ~ | ~ |
| Body substance sample | Body substance sample | Body substance sample (63) | 498 | \checkmark | ~ | \checkmark | ~ | ~ | ~ | \checkmark |
| Fluid sample | Fluid sample | Fluid sample (50) | 257 | \checkmark | ~ | \checkmark | ~ | ~ | ~ | \checkmark |
| Bone marrow specimen | Bone marrow specimen | Bone marrow specimen (8) | | \checkmark | ~ | \checkmark | _ | _ | _ | _ |
| Bone specimen | Specimen from bone | Musculoskeletal sample (15) | | - | - | - | - | - | - | - |
| Specimen from nervous system | Specimen from nervous system | Specimen from nervous system (12) | | \checkmark | \checkmark | ~ | \checkmark | \checkmark | \checkmark | \checkmark |
| Dermatological specimen | Dermatological sample | Dermatological sample (8) | | \checkmark | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| Device specimen | Device specimen | Device specimen (19) | | \checkmark | \checkmark | ~ | \checkmark | \checkmark | \checkmark | \checkmark |
| Digestive system specimen | Specimen from digestive system | Specimen from digestive system (50) | | \checkmark | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| Endocrine system specimen | Endocrine sample | Endocrine sample (10) | 26 | \checkmark | ✓ | ✓ | ~ | ~ | ~ | _ |
| Genital system specimen, male | Male genital sample | Specimen from trunk (132) | 489 | _ | _ | _ | _ | _ | _ | _ |
| Genitourinary specimen | Genitourinary sample | Specimen from trunk (132) | 489 | _ | _ | _ | _ | _ | _ | _ |
| Hair specimen, scalp | Hair specimen | Dermatological sample (8) | 30 | _ | _ | _ | _ | _ | _ | _ |
| Musculoskeletal specimen | Musculoskeletal sample | Musculoskeletal sample (15) | 56 | \checkmark | \checkmark | ~ | \checkmark | \checkmark | \checkmark | \checkmark |
| Skin specimen | Specimen from skin | Dermatological sample (8) | 30 | _ | _ | _ | _ | _ | _ | _ |
| Soft tissue specimen | Soft tissue sample | Soft tissue sample (21) | 92 | \checkmark | ✓ | ✓ | ✓ | ~ | ✓ | \checkmark |
| Cardiovascular sample | Cardiovascular sample | Cardiovascular sample (12) | 28 | \checkmark | ✓ | ✓ | ✓ | ~ | ✓ | _ |
| Specimen from eye | Specimen from eye | Specimen from head and neck structure (53) | 196 | _ | _ | _ | _ | _ | _ | _ |
| Specimen from joint | Joint sample | Musculoskeletal sample (15) | 56 | _ | _ | _ | _ | _ | _ | _ |
| Lesion sample | Lesion sample | Lesion sample (17) | 118 | \checkmark | ✓ | ✓ | ✓ | ~ | ✓ | \checkmark |
| Stool specimen | Stool specimen | Body substance sample (63) | 498 | _ | _ | _ | _ | _ | _ | _ |
| # Identified topics (C) | - | | | 13 | 13 | 13 | 12 | 12 | 12 | 10 |

Table 1– Identification results for 21 chosen topics in weighted aggregate taxonomies with different thresholds b

Major Topic Identification

We experiment iterating the threshold b over the range 1...30 and generate the weighted aggregate taxonomy for each b. Each such weighted aggregate taxonomy is inspected to determine its effectiveness in capturing major topics. Precision, recall, and F measure [9] were calculated for each weighted aggregate taxonomy, with the expert's topic list serving as a gold standard.

As a preliminary experiment, we tested how many of the gold standard topics appeared as partial-areas in the taxonomy (not the weighted aggregate taxonomies). We found that, out of the 21 topics chosen by the expert, 13 appear as partial-areas. This yields a recall of 0.62 (13/21) and, with 503 partial-areas in the taxonomy, very low precision of 0.03 (13/503). Note that many partial-areas are very small. In contrast, the weighted aggregate taxonomy, which eliminates the small partial-areas, is more effective. To balance recall and precision, we chose the weighted aggregate taxonomy with the *b* value that maximizes the F measure.

Results

If the root concept r of a partial-area appears in the weighted aggregate taxonomy of threshold b, then r is considered a topic identified by that weighted aggregate taxonomy, a corresponding checkmark "√" is placed in Table 1. Otherwise, a dash "-" is written. For example, the topic Bone marrow specimen is captured by a partial-area Bone marrow specimen (8) with an aggregated weight 13 (Table 1). Therefore, it is identified by all weighted aggregate taxonomies with $b \leq 13$ (b=1, 5, 10) However, for b>13, Bone marrow specimen (8) is folded into an ancestor partial-area and disappears. No weighted aggregate taxonomy with b > 13 identifies the topic Bone marrow specimen. As another example, Bone specimen was not identified by the weighted aggregate taxonomy with any b value as major topic (Row 5 of Table 1), since its mapped SNOMED CT concept Specimen from bone (Row 5, Column 2 of Table 1) is not a root of a partial-area.

At the bottom of Table 1, we show the totals of the identified

topics for the respective taxonomies. For example, for b=5, the total is 13. Table 2 shows each weighted aggregate taxonomy's number of partial-areas (A), recall, precision, and F. Recall is the ratio of identified topics and total topics (R=C/S, where S=21). Precision is the ratio of the identified topics and the number of partial-areas (P=C/A). For example, for b=25, the number of partial-areas is 29, the number of identified topics is 12, R=0.57, P=0.41 & F=0.48. Table 2 shows that b=25 yields the taxonomy where F is maximized. In this case, the weighted aggregate taxonomy captures 12 of the 21 topics. Figure 4 shows this weighted aggregate taxonomy with the 12 partial-areas identifying topics highlighted in yellow. The total number of concepts in these 12 aggregate partial-areas is 988, accounting for 61.0% (988/1620) of the concepts in the *Specimen* hierarchy.

Table 2–Parameterize performance of weighted aggregate taxonomies for various thresholds

| | - | | | | | | |
|-------------------------------|------|------|------|------|------|------|------|
| <i>b</i> = | 1 | 5 | 10 | 15 | 20 | 25 | 30 |
| # Identified topics | 13 | 13 | 13 | 12 | 12 | 12 | 10 |
| # Partial-areas (A) | 503 | 89 | 54 | 40 | 35 | 29 | 26 |
| Recall $(R = C/S)$ | 0.62 | 0.62 | 0.62 | 0.57 | 0.57 | 0.57 | 0.48 |
| Precision $(P = C/A)$ | 0.03 | 0.15 | 0.24 | 0.30 | 0.34 | 0.41 | 0.38 |
| $F = 2 \cdot P \cdot R/(P+R)$ | 0.05 | 0.24 | 0.35 | 0.39 | 0.43 | 0.48 | 0.43 |

As an ancillary experiment, we carried out a feedback step with our domain expert (GE). When inspecting the weighted aggregate taxonomy for threshold b, one can assess whether its other partial-areas beyond those in the gold standard list are worthy of the designation "major topic", for example, those aggregate partial-areas (Figure 4) categorizing over 25 concepts, but not in the given list. Since some topics may have been overlooked originally due to various reasons, e.g. Specimen from head and neck structure, a compound topic name with two body parts, and Tissue specimen obtained by excision, corresponding to two relationships Specimen procedure and Specimen source topography. Figure 4 was shown to (GE). He manually determined that 13 more partialareas, highlighted in pink, warranted inclusion in the list of major specimen topics, while the other three (in white) are deemed as non-major topics. Reevaluating the experiment (with 21+13=34 major topics), we obtained R=0.74



Figure 4-Weighted aggregate taxonomy for the Specimen hierarchy with b=25. The 12 partial-areas corresponding to the original given topics are highlighted in yellow. The 13 topics added during the enhancement step are highlighted in pink

(=25/34), P=0.86 (=25/29) and F=0.79 for b=25. The number of concepts in these 25 aggregate partial-areas is 1,524 (94.1%) of the concepts in the *Specimen* hierarchy).

Discussion

Summarizing a large ontology is a challenge as there is a lack of an objective universally accepted criteria for what constitutes a "good summarization" of an ontology. Various applications require different summaries of various granularites. Nevertheless, the management of ontologies requires "big picture" comprehension that can be enabled by compact summarization networks such as *weighted aggregate taxonomies*. Our technique is applicable to any ontology. We demonstrated our technique using SNOMED CT because of its importance in clinical applications and its large size.

We used a knowledge-oriented approach, where the importance of a topic is based on the number of concepts related to that topic in an ontology. To measure the quality of our summarizations we compared the number of identified major topics with a gold standard list of topics selected by a domain expert, who selected topics from a clinical perspective. The performance of our technique was optimized by maximizing the F symmetric measure. We modified partial-area taxonomies into weighted aggregate taxonomies to overcome difficulties in bridging the gap between the clinical perspective of the given gold standard list and the knowledge better results.

A future problem is accounting for major topics that are summarized by large partial-areas (e.g., *Male genital sample* is in *Specimen from trunk* (132)). Such a topic is missed by a taxonomy since it does not appear as a root of a partial-area. It would also be more objective to obtain the gold standard list from another ontology (or other authorative sources) rather than from one of the authors (GE). Thus, we plan to experiment with the *Disease, Disorder or Finding* hierarchy of NCIt, using ICD-10 as a source of the topic list. We are also planning evaluation studies to compare our technique to related summarization methods, such as key-concept-based ontology summarization [10] and information-content-based approaches [11].

Conclusions

Summarizing knolwedge bases for maintenance and "big picture" comprehension is a Big Knowledge challenge. Taxonomies were created to summarize structure and content of ontologies. We modify taxonomies to properly support "big picture" comprehension. For this, we presented the *weighted aggregate partial-area taxonomy* and

measured its performance for "Big Knowledge" coverage of a domain expert's gold standard list of major topics for SNOMED CT's *Specimen* hierarchy. Our methodology was parameterized in terms of a threshold b for node size to obtain optimal performance by maximizing the symmetric F measure balancing recall and precision. A manual enhancement step improved the algorithmic results.

Acknowledgments

Research reported here was supported by the National Cancer Institute of the National Institutes of Health under Award Number R01CA190779. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

References

- V.L. Patel, D.R. Kaufman, and J. Arocha, Conceptual change in the biomedical and health sciences domain, *Advances in instructional* psychology 5 (2000), 329–92.
- [2] K.W. Fung, R. Richesson, and O. Bodenreider, Coverage of rare disease names in standard terminologies and implications for patients, providers, and research, AMIA Symp (2014), 564–72.
- [3] C. Ochs, Y. Perl, J. Geller, and M.A. Musen, Using aggregate taxonomies to summarize SNOMED CT evolution, in: *BIBM Conference*, 2015, pp. 1008–15.
- [4] Y. Wang, M. Halper, H. Min, Y. Perl, Y. Chen, and K.A. Spackman, Structural methodologies for auditing SNOMED, *J Biomed Inform* 40 (2007), 561–81.
- [5] M.Q. Stearns, C. Price, K.A. Spackman, and A.Y. Wang, SNOMED clinical terms: overview of the development process and project status, *AMIA Symp* (2001), 662–66.
- [6] K. Giannangelo and S.H. Fenton, SNOMED CT survey: an assessment of implementation in EMR/EHR applications, *Perspect Health Inf Manag* 5 (2008), 7.
- [7] D.L. Rubin, N.H. Shah, and N.F. Noy, Biomedical ontologies: a functional perspective, *Brief Bioinform* 9 (2008), 75–90.
- [8] N. Li, E. Motta, and M. d'Aquin, Ontology summarization: an analysis and an evaluation, in: *IWEST*, 2010.
- [9] C.D. Manning, P. Raghavan, and H. Schütze, Evaluation in information retrieval, in: *Introduction to Information Retrieval*, Cambridge University Press, 2008, pp. 151–75.
- [10] S. Peroni, E. Motta, and M. d'Aquin, Identifying key concepts in an ontology, through the integration of cognitive principles with statistical and topological measures, in: ASWC, 2008, 242–56.
- [11] G.K. Mazandu and N. J. Mulder, Information Content-Based Gene Ontology Functional Similarity Measures: Which One to Use for a Given Biological Data Type? *PLoS One* 9 (2014), e113859.

Address for Correspondence

Ling Zheng, <u>Iz265@njit.edu</u>, NJIT, University Heights, Newark, NJ 07102-1982 USA