

3

Lexical, Terminological, and Ontological Resources for Biological Text Mining

Olivier Bodenreider

3.1 Introduction

Biomedical terminologies and ontologies are frequently described as enabling resources in text mining systems [1–3]. These resources are used to support tasks such as entity recognition (i.e., the identification of biomedical entities in text), and relation extraction (i.e., the identification of relationships among biomedical entities). Although a significant part of current text mining efforts focuses on the analysis of documents related to molecular biology, the use of lexical, terminological, and ontological resources is mentioned in research systems developed for the analysis of clinical narratives (e.g., MedSyndikate [4]), or the biological literature (e.g., BioRAT [5], GeneScene [6], EMPathIE [7], and PASTA [7]). Of note, some systems initially developed for extracting clinical information later have been adapted to extract relations among biological entities (e.g., MedLEE [8], GENIES [9], and SemRep/SemGen [10]). Commercial systems such as TeSSI, from Language & Computing (<http://www.landcglobal.com/>) also make use of such resources.

Entity recognition often draws on lists of entity names collected in lexicons, gazetteers, and, more generally, terminology resources. For example, lists of disease names can be easily extracted from disease resources, such as the International Classification of Diseases (ICD); from the disease component of general resources, such as the Medical Subject Headings (MeSH); and from specialized resources, such as the Online Multiple Congenital Anomaly/Mental

Retardation (MCA/MR) Syndromes. In contrast, relation extraction may benefit from the relationships represented among terms in terminologies (e.g., ‘Parkinson’s disease’ *child of* ‘Neurodegenerative diseases’ in MeSH), and in ontologies [e.g., ‘Basal ganglia’ *finding site of* ‘Parkinson’s disease’ in the Systematized Nomenclature of Medicine Clinical Terms (SNOMED CT)].

Biomedical lexicons, such as the UMLS Specialist lexicon, collect lexical items (i.e., words and multiword expressions) frequently observed in biomedical text corpora, and record information about them, including parts of speech (e.g., noun or adjective), inflectional variants (e.g., singular or plural), and spelling variants (e.g., American versus British English). This information is useful not only to NLP tools, such as part-of-speech taggers and parsers, but also to entity recognition systems, since it can help identify variants of entity names in text [11].

The purpose of biomedical terminology is to collect the names of entities employed in the biomedical domain. Most biomedical terminologies record synonymous terms (e.g., *Parkinson’s disease* and *Paralysis agitans*), and have some kind of hierarchical organization, often treelike or graphlike [12]. Terminology-driven approaches to text mining have been explored in [13].

In contrast, biomedical ontology aims to study entities (i.e., substances, qualities, and processes) of biomedical significance, and the relations among them. Examples of such entities include substances such as the mitral valve and glucose, qualities such as the diameter of the left ventricle and the catalytic function of enzymes, and processes such as blood circulation and secreting hormones. Fundamental relations in biomedical ontologies include not only *is a* and *part of*, but also *instance of*, *adjacent to*, *derives from*, and so forth [14].

In practice, the distinction between lexicons, terminologies, and ontologies is not always sharp. On the one hand, although ontologies mostly focus on relations among entities, some of them also record the names by which entities are referred. On the other hand, although terminologies essentially collect the names of entities, their hierarchical organization also reflects relations among such entities. Finally, the very names of these resources can be misleading. For example, despite its name, the GO defines itself as a controlled vocabulary (i.e., a terminological resource), but like ontologies, its terms are linked by relationships such as *is a* and *part of*. However, the definition and use of such relations is not consistent throughout GO [15], as would be expected from ontologies.

The objective of this chapter is to present some of the resources (lexicons, terminologies, and ontologies) of interest for entity recognition and relation extraction tasks. Providing an exhaustive list of these resources is beyond the scope of this chapter. Many of these resources are highly specialized, and therefore would be of little interest to most readers. Instead, we have selected general,

publicly available resources that have been shown to be useful for biomedical text mining. This review is purposely limited to resources in English.

We start by presenting an extended example, illustrating biomedical terms in two pieces of text. We then give a brief description of the major resources available, with a particular emphasis on the UMLS [16]. Finally, we discuss some issues related to biomedical terms and biomedical relations. The reader is referred to Chapters 6 and 7 for a detailed presentation of the tasks of entity recognition and relation extraction.

3.2 Extended Example

In this example, we consider two short pieces of text related to the genetic disease neurofibromatosis 2. This is an autosomal dominant disease, characterized by tumors called schwannomas that involve the acoustic nerve, as well as other features [17]. The disorder is caused by mutations of the *NF2* gene, resulting in absence or inactivation of the protein product. The protein product of *NF2* is commonly called merlin (but also neurofibromin 2 and schwannomin), and functions as a tumor suppressor. The first fragment of text (3.1) is extracted from the abstract of an article [18]. The second is the definition of neurofibromatosis 2 in the MeSH vocabulary (<http://www.nlm.nih.gov/mesh/>).

- (3.1) Neurofibromatosis type 2 (NF2) is often not recognized as a distinct entity from peripheral neurofibromatosis. NF2 is a predominantly [intracranial condition] whose hallmark is [bilateral vestibular schwannomas]. NF2 results from a mutation in the gene named merlin, located on chromosome 22.
- (3.2) Neurofibromatosis type 2: An [autosomal dominant disorder] characterized by a high incidence of [bilateral acoustic neuromas] as well as schwannomas of other [cranial and peripheral nerves], and other [benign intracranial tumors] including meningiomas, ependymomas, spinal neurofibromas, and gliomas. The disease has been linked to mutations of the NF2 gene on chromosome 22 (22q12) and usually presents clinically in the first or second decade of life.

3.2.1 Entity Recognition

Many biomedical entities can be identified in these two fragments. Underlined expressions correspond to terms present in the UMLS Metathesaurus. This is the case, for example, of the disease *neurofibromatosis 2* and the protein *merlin*. Interestingly, *vestibular schwannomas* in (3.1) and *acoustic neuromas* in (3.2), although lexically distinct, name the same tumor. While a lexicon is useful to

identify these disease names, a terminology (or ontology) is required to identify them as synonymous. These two terms are names for the same disease concept in the UMLS Metathesaurus (C0027859). The list of UMLS concepts that can be identified in the two text fragments is given in Table 3.1.

Many expressions extracted from the two text fragments can be mapped to the UMLS Metathesaurus through a simple match (i.e., exact match or after normalization). Except for *merlin*, which maps to both a protein and a bird, the mapping is unambiguous. In contrast, expressions of (3.1) and (3.2) enclosed in brackets also correspond to biomedical entities, but the name found in the text cannot be mapped directly to a UMLS concept. Expressions, such as *intracranial condition* in (3.1), are vague, compared to the corresponding concept names in the UMLS (e.g., ‘central nervous system diseases’). Complex phrases, such as *cranial and peripheral nerves* in (3.2), refer to two concepts (i.e., ‘cranial nerves’ and ‘peripheral nerves’) present in the Metathesaurus. Conversely, some expressions in the text convey more precision than the corresponding concepts found in biomedical terminologies [e.g., *bilateral vestibular schwannomas* in (3.1) versus ‘vestibular schwannomas’ and *benign intracranial tumors* in (3.2) versus ‘intracranial tumors’]. In these cases, while terminological resources are useful for identifying entities in text, they may not be sufficient for capturing all nuances present in the text. Term variation and management issues are discussed extensively in Chapter 4.

3.2.2 Relation Extraction

Once entities have been identified in text fragments, the next step consists of identifying the relationships among them, such as *vestibular schwannomas manifestation of neurofibromatosis 2* and *NF2 gene located on chromosome 22*. Such relations may be explicitly represented in biomedical ontologies. For example, the relation ‘schwannomas’ *associated morphology of* ‘neurofibromatosis 2’ is asserted in SNOMED CT. However, ontologies do not necessarily contain such fine-grained assertions, but may rather represent higher-level facts such as ‘gene’ *located on* ‘chromosome’. A relation extraction system would first identify *NF2 gene* as a kind of gene and *chromosome 22* as a kind of chromosome, before inferring that a particular gene (*NF2 gene*) is located on a particular chromosome (*chromosome 22*).

The use of ontologies to support relation extraction often requires the system to identify in the text not only entities, but also potential relationships. Clues for identifying relationships include lexical items (e.g., the preposition ‘on’ for the relationship *located on*), and syntactic structures (e.g., *intracranial tumors including meningiomas* for *meningiomas is a intracranial tumors*), as well as statistical and pattern-based clues (not presented here). Relations may span several sentences, and their identification often requires advanced linguistic

Table 3.1
UMLS Concepts: Identifiers, Preferred Names, and Semantic Types, Identifiable in Text Fragments (3.1) and (3.2)

Source	String in Text	M	CUI	Preferred Name	S. Types
(1) (2)	Neurofibromatosis type 2	s	C0027832	Neurofibromatosis 2	neop
(1)	NF2	s	C0085114	Neurofibromatosis 2 genes	gngm
(1)	peripheral neurofibromatosis	s	C0027831	Neurofibromatosis 1	neop
(1)	{intracranial condition}	—	C0007682	Central Nervous System Diseases	dsyn
(1)	{bilateral} vestibular schwannomas	a	C0027859	Neuroma, Acoustic	neop
(1) (2)	mutation/mutations	s	C0026882	Mutation	genf
(1)	gene	s	C0017337	Genes	gngm
(1)	merlin	m	C0254123	Neurofibromin 2	aapp, bats
(1) (2)	chromosome 22	s	C0008665	Chromosomes, Human, Pair 22	celc
(2)	autosomal dominant disorder	a	C0265385	Autosomal dominant hereditary disorder	dsyn
(2)	bilateral acoustic neuromas	s	C1136042	Neuroma, Acoustic, Bilateral	neop
(2)	schwannomas	s	C0027809	Neurilemmoma	neop
(2)	cranial and peripheral nerves	—	C0010268	Cranial Nerves	bpoc
(2)	{benign} intracranial tumors	a	C0031119	Peripheral Nerves	bdsy
(2)	meningiomas	s	C0750978	Neoplasms, Intracranial	neop
(2)	ependymomas	s	C0025286	Meningioma	neop
(2)	neurofibromas	s	C0014474	Ependymoma	neop
(2)	gliomas	s	C0027830	Neurofibroma	neop
(2)	disease	s	C0017638	Glioma	neop
(2)	NF2 gene	s	C0012634	Disease	dsyn
(2)		s	C0085114	Neurofibromatosis 2 genes	gngm

CUI: identifier (see Table 3.4 for the full names).

Column 'M' indicates the type of match (s: single simple match, m: multiple simple matches, a: approximate match, -: no direct match).

techniques, such as anaphora and coreference resolution. For example, from the last sentence of (3.2), the relation *disease associated with mutation* can be extracted. While accurate, this relation is incomplete in this context, because *disease* actually refers not to any disease, but to *neurofibromatosis 2* (anaphoric relation). Similarly, *mutations of the NF2 gene* (not mutations in general) is the entity associated with the disease. Therefore, the complete relation to be extracted is *neurofibromatosis 2 associated with mutations of the NF2 gene*. The potential relations extracted from the text then can be validated against the relations explicitly represented in the ontology, or inferred from it.

3.3 Lexical Resources

The resources presented under this category provide the lexical and lexico-syntactic information needed for parsing text. The major resource for biomedical text is the Specialist lexicon. Additionally, specialized resources can be useful for analyzing subdomains of biomedicine (e.g., lists of gene names for molecular biology corpora). Conversely, general resources such as WordNet also can help analyze the literature written for less-specialized audiences (e.g., patients).

3.3.1 WordNet

WordNet is an electronic lexical database developed at Princeton University, which serves as a resource for applications in natural language processing and information retrieval [19]. The core structure in WordNet is a set of synonyms (synset) that represents one underlying concept. For example, the synset representing ‘hemoglobin’ also contains the lexical entries ‘haemoglobin’ (British-English spelling) and ‘Hb’ (abbreviation). A definition is provided for the synset: “a hemoprotein composed of globin and heme that gives red blood cells their characteristic color; function primarily to transport oxygen from the lungs to the body tissues.” There are separate structures for each linguistic category covered: nouns, verbs, adjectives, and adverbs. For example, the adjective ‘renal’ and the noun ‘kidney’, although similar in meaning, belong to two distinct structures, and a specific relationship (pertainymy) relates the two forms. The current version of WordNet (2.0) contains over 114,000 noun synsets. In addition to being a lexical resource, WordNet has some of the features of an ontology. For example, each synset in the noun hierarchy belongs to at least one *is a* tree (e.g., ‘hemoglobin’ *is a* ‘protein’), and may additionally belong to several *part of*-like trees (‘hemoglobin’ *substance of* ‘red blood cell’). Because of its modest coverage of the biomedical domain [20, 21], WordNet has been used only in a limited number of projects in biomedicine [22], where resources such as the UMLS usually play a more prominent role. WordNet is available free of charge

from <http://wordnet.princeton.edu/>. Application programming interfaces (API) have been developed for the major programming languages, making it relatively easy for developers to integrate WordNet into their applications.

3.3.2 UMLS Specialist Lexicon

The Specialist lexicon is one of three knowledge sources developed by the NLM, as part of the UMLS project. It provides the lexical information needed for processing natural language in the biomedical domain [23]. The lexicon entry for each word or multiword term records syntactic (part-of-speech, allowable complementation patterns), morphological (base form, inflectional variants), and orthographic (spelling variants) information. It is, in fact, a general English lexicon that includes many biomedical terms. Lexical items are selected from a variety of sources, including lexical items from MEDLINE/PubMed citation records, the UMLS Metathesaurus, and a large set of lexical items from medical and general English dictionaries. Contrary to WordNet, the Specialist lexicon does not include any information about synonymy or semantic relations among its entries. However, this information is present in the Metathesaurus, another component of the UMLS (see Section 3.4.3). The record for ‘hemoglobin’ in the Specialist lexicon, shown in Figure 3.1, indicates the base form, one spelling variant, and two inflectional classes, since hemoglobin is used as both a mass noun (e.g., in *Hemoglobin concentration is reported as grams of hemoglobin per deciliter of blood*), and as a countable (e.g., in *the study of hemoglobins, both normal and mutant*). Additionally, the abbreviation ‘Hb’ and the acronym ‘Hgb’ are cross-referenced to ‘hemoglobin’. The Specialist lexicon is distributed as part of the UMLS, and can be queried through application programming interfaces for Java and XML. It is also available as an open source resource, as part of the Specialist NLP tools (<http://SPECIALIST.nlm.nih.gov>).

3.3.3 Other Specialized Resources

While general resources such as WordNet and the Specialist lexicon provide a good coverage of the general biomedical language, they (purposely) fail to cover

```

{
  base=hemoglobin           (base form)
  spelling_variant=haemoglobin
  entry=E0031208           (identifier)
  cat=noun                 (part of speech)
  variants=uncount        (no plural)
  variants=reg            (plural: hemoglobins, haemoglobins )
}

```

Figure 3.1 Representation of ‘hemoglobin’ in the Specialist lexicon.

in detail specialized subdomains, such as gene and protein names, or chemical and drug names. Therefore, the syntactic analyzers and parsers relying on these resources may give suboptimal results when analyzing specialized corpora (e.g., molecular biology abstracts). One approach to solving this problem is to use machine learning techniques to identify the names of specialized entities. Alternatively, or in conjunction with these techniques, resources, such as lists of gene, protein, chemical, and drug names, can be exploited [24]. In molecular biology, for example, the Human Genome Organization (HUGO) has established, through its Gene Nomenclature Committee (HGNC), a list of over 20,000 approved gene names and symbols, called Genew [25]. Recorded in this database are the symbol 'NF2' and the name 'neurofibromin 2 (bilateral acoustic neuroma)' for the gene merlin, whose mutation causes the disease neurofibromatosis 2. More generally, lists of names for specialized entities can be extracted from specialized resources. Examples of publicly available specialized resources for genes, proteins, chemical entities, and drugs are given in Table 3.2. Finally, acronyms and abbreviations harvested from the biomedical literature [26, 27] and collected in databases [28] also can benefit entity recognition applications. This issue is discussed extensively in Chapter 6.

3.4 Terminological Resources

The purpose of terminology is to collect the names of entities employed in the biomedical domain [29]. Terminologies typically provide lists of synonyms for the entities in a given subdomain and for a given purpose. As such, they play an important role in entity recognition. Additionally, most terminologies have some kind of hierarchical organization that can be exploited for relation extraction purposes. Many terminologies consist of a tree, where nodes are terms, and links represent parent-to-child or more-general-to-more-specific relationships. Some terminologies allow multiple inheritance, and have the structure of a directed acyclic graph. The Gene Ontology and MeSH provide examples of terminological systems created to support different tasks. Because it integrates a large number of terminologies, the UMLS Metathesaurus is the terminological system most frequently used in the analysis of biomedical texts.

3.4.1 Gene Ontology

The Gene Ontology (GO) is a controlled vocabulary, developed by the Gene Ontology Consortium, for the annotation of gene products in model organisms. GO is organized in three separate hierarchies—for molecular functions (6,933 terms), biological processes (9,053 terms), and cellular components (1,414 terms), as of February 1, 2005 [30]. For example, annotations for the

Table 3.2

Examples of Publicly Available Specialized Resources for Genes, Proteins, Chemical Entities, and Drugs

Domain	Resources	URL
Genes and proteins	Genew	http://www.gene.ucl.ac.uk/nomenclature/
	Entrez Gene (formerly LocusLink)	http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?db=gene
	UniProt	http://www.ebi.uniprot.org/index.shtml
Chemical entities	PubChem	http://pubchem.ncbi.nlm.nih.gov/
	ChemIDplus	http://chem.sis.nlm.nih.gov/chemidplus/chemidlite.jsp
	ChEBI	http://www.ebi.ac.uk/chebi/
Drugs	RxNorm	http://www.nlm.nih.gov/research/umls/rxnorm_main.html
	National Drug Code	http://www.fda.gov/cder/ndc/

gene *NF2* in the Gene Ontology Annotation (GOA) database (<http://www.ebi.ac.uk/GOA/>) include the molecular function term ‘cytoskeletal protein binding’, the biological process term ‘negative regulation of cell proliferation’, and the cellular component terms ‘plasma membrane’ and ‘cytoskeleton’. Each of the three hierarchies is organized in a directed acyclic graph, in which the nodes are GO terms, and the edges represent the GO relationships *is a* and *part of*. For example, as illustrated in Figure 3.2, the relations of the cellular component ‘cytoskeleton’ to its parent terms include ‘cytoskeleton’ *is a* ‘intracellular nonmembrane-bound organelle’ and ‘cytoskeleton’ *part of* ‘intracellular’. GO terms may have synonyms (e.g., synonyms for ‘plasma membrane’ include ‘cytoplasmic membrane’ and ‘plasmalemma’). Most terms have a textual definition (e.g., for ‘plasma membrane’: “The membrane surrounding a cell that separates the cell from its external environment. It consists of a phospholipid bilayer and associated proteins.”).

Both the names and the relations comprised in the Gene Ontology can benefit text mining applications. The names of molecular functions, biological processes, and cellular components are frequently used in the biomedical literature [31]. For example, the biological process ‘activation of MAPK’ and the cellular component ‘adherens junction’ can be identified in the title *Erbin regulates MAP kinase activation and MAP kinase-dependent interactions between merlin and adherens junction protein complexes in Schwann cells*. As illustrated in the

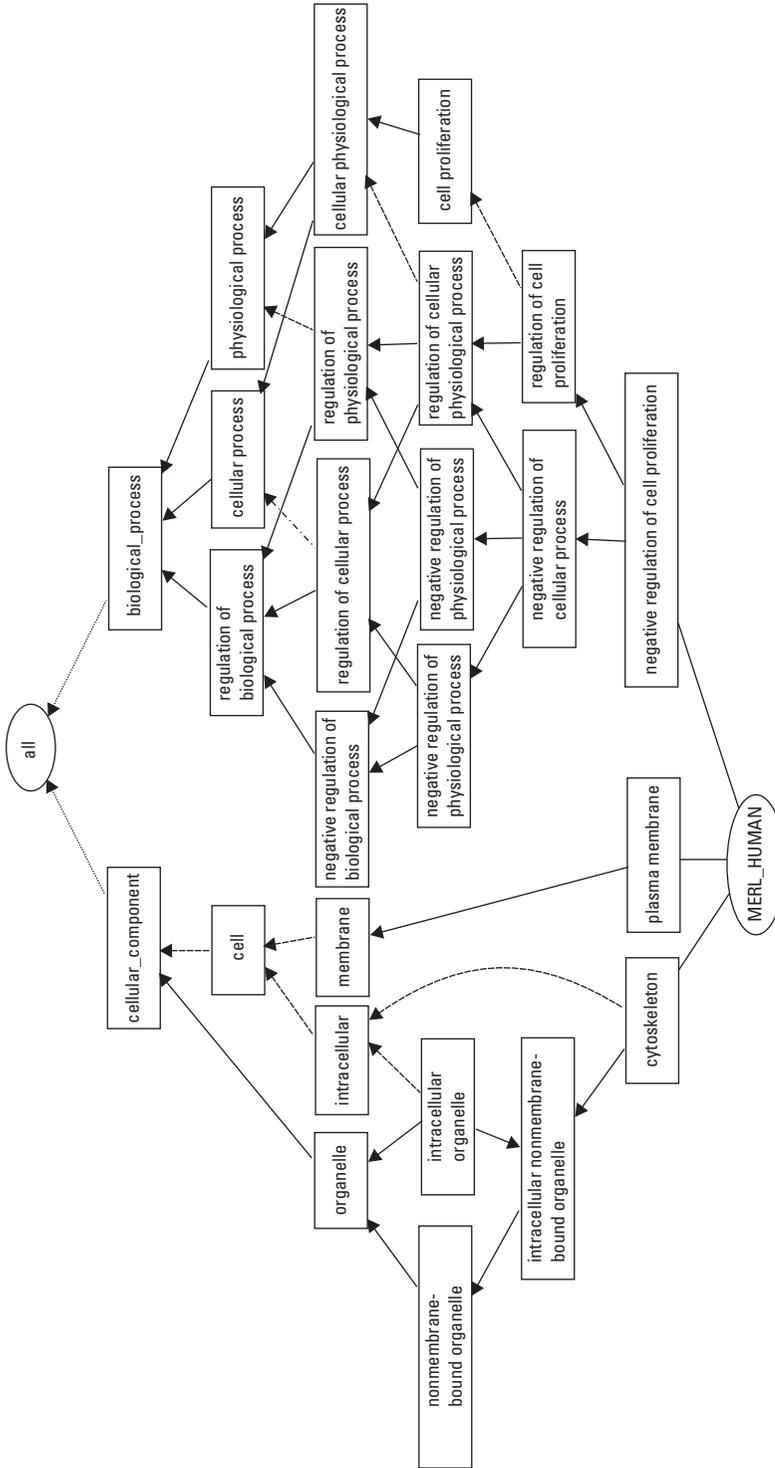


Figure 3.2 Representation of the gene product *merlin* ('MEL_HUMAN'). Solid lines represent some of its annotations in the Gene Ontology. Solid and dashed arrows represent the Gene Ontology relationships *is a* and *part of*, respectively.

following text fragment (3.3), hierarchical relations can help resolve anaphora and interpret associative relations.

- (3.3) The organization of the actin cytoskeleton in prefusion aligning myoblasts is likely to be important for their shape and interaction. We investigated actin filament organization and polarity by transmission electron microscopy (TEM) in these cells.

The terms *actin cytoskeleton* and *actin filament* identified in the first two sentences of (3.3) are present in GO. Moreover, a relation between them is explicitly recorded in GO ('actin filament' *part of* 'actin cytoskeleton'), which helps link together the two sentences. However, many concepts and relations are not represented in GO, or other biomedical terminologies. For example, a relation between *myoblasts* and *these cells*—namely, 'myoblast' *is a* 'cell'—is needed to resolve the anaphoric relation between the two terms in (3.3). Such a relation cannot be found in GO, where the term *myoblast* is not even represented.

Finally, GO terms constitute an entry point to annotation databases, providing a wealth of relations between gene products and the molecular functions, biological processes, and cellular components with which they are associated (e.g., 'NF2' *has biological process* 'negative regulation of cell proliferation'). GO is available from <http://geneontology.org/>, and is distributed in various formats, including XML and database formats. Perl and Java application programming interfaces are also available. GO is one of the source vocabularies included in the UMLS Metathesaurus. GO is a member of a family of controlled vocabularies, called Open Biomedical Ontologies (OBO). These resources can be useful in text mining applications as a source of specialized vocabulary (e.g., for chemicals or experimental conditions). OBO resources are available at <http://obo.sourceforge.net>.

3.4.2 Medical Subject Headings

The Medical Subject Headings (MeSH) thesaurus is a controlled vocabulary, produced by the National Library of Medicine, and used for indexing, cataloging, and searching for biomedical and health-related information and documents [32]. It consists of 22,995 descriptors (main headings), organized in 15 hierarchies. Additionally, a set of approximately 150,000 "supplementary concept records" provides a finer-grained representation of biomedical entities, including chemicals and proteins. A list of entry terms (synonyms or closely related terms) is given for each descriptor. Entry terms for the disease 'Neurofibromatosis 2' include *Neurofibromatosis Type II*, *Bilateral Acoustic Neurofibromatosis*, *Bilateral Acoustic Schwannoma*, and *Familial Acoustic Neuromas*. A scope note often provides a definition of the descriptor. In the

MeSH thesaurus, descriptors are related by parent/child relations; each descriptor has at least one parent, and may have several. For example, ‘Neurofibromatoses’ and ‘Neuroma, Acoustic’ are the two parents of the descriptor ‘Neurofibromatosis 2’. The arrangement of MeSH descriptors in hierarchies is intended to serve the purpose of indexing and information retrieval, and does not always follow strict classifications. In addition to hierarchical relations, cross references may link a descriptor to descriptors from other hierarchies. For example, the disease ‘Neurofibromatosis 2’ is linked to the protein ‘Neurofibromin 2’ and to the gene ‘Genes, Neurofibromatosis 2’. The MeSH thesaurus is used by the NLM for indexing articles from 4,600 biomedical journals for the MEDLINE/PubMed database. Like GO, MeSH can be used in text mining applications, due to the many names and relations it provides. Its scope is broader than that of GO, but its granularity is coarser. MeSH is available from <http://www.nlm.nih.gov/mesh/> in various formats, including XML. MeSH is one of the source vocabularies included in the UMLS Metathesaurus.

3.4.3 UMLS Metathesaurus

The UMLS Metathesaurus is one of three knowledge sources developed and distributed by the NLM, as part of the UMLS project [16]. Version 2005AA of the Metathesaurus contains over 1 million biomedical concepts and 5 million concept names, from more than 100 controlled vocabularies and classifications (some in multiple languages) used in patient records, administrative health data, bibliographic and full-text databases, and expert systems. The Metathesaurus also records over 16 million relations among these concepts, either inherited from the source vocabularies or specifically generated. While the Metathesaurus preserves the names, meanings, hierarchical contexts, attributes, and interterm relationships present in its source vocabularies, it also integrates existing terminologies into a common semantic space. As in WordNet, synonymous names are clustered together to form a concept. The Metathesaurus also assigns a unique identifier to each concept, and establishes new relations between terms from different source vocabularies, as appropriate. Each concept is also categorized with at least one semantic type from the UMLS Semantic Network (see Section 3.5.2), independent of its hierarchical position in the source vocabularies. The scope of the Metathesaurus is determined by the combined scope of its source vocabularies, including Gene Ontology and MeSH, disease vocabularies (e.g., International Classification of Diseases), clinical vocabularies (e.g., SNOMED CT), nomenclatures of drugs and medical devices, as well as the vocabularies of many subdomains of biomedicine (e.g., nursing, psychiatry, and gastrointestinal endoscopy).

Examples of Metathesaurus concepts are given in Table 3.1. C0254123 identifies the protein ‘neurofibromin 2’, whose synonyms include *merlin*, *NF2*

protein, and *schwannomin*. Its semantic types are ‘Amino Acid, Peptide, or Protein’ and ‘Biologically Active Substance’. The following source vocabularies contributed names to this concept: MeSH, SNOMED CT, and the NCI Thesaurus. Once integrated in the Metathesaurus, ‘neurofibromin 2’ has multiple parents including ‘membrane proteins’ (from MeSH), ‘tumor suppressor proteins’ (from both MeSH and SNOMED CT), and ‘signaling protein’ (from the NCI Thesaurus). Its only descendant is ‘merlin, *Drosophila*’ (from MeSH). Beside hierarchical relations, associative relations link the protein ‘neurofibromin 2’ to the gene ‘neurofibromatosis 2 genes’ and to the disease ‘neurofibromatosis 2’. The frequencies of co-occurrence of MeSH descriptors in MEDLINE/PubMed citations are also recorded in the Metathesaurus. For example, during the last 10 years, the descriptors ‘Neurofibromin 2’ and ‘Neurofibromatosis 2’ occurred together 13 times as major descriptors. The descriptors ‘Membrane Proteins’ (8 times), ‘Phosphoproteins’ and ‘NF2 gene’ (7 times), and ‘Cell Transformation, Neoplastic’ (5 times) frequently co-occur with ‘Neurofibromin 2’.

Section 3.2.1 illustrated how the Metathesaurus can be used in entity recognition and relation extraction tasks. Used in many biomedical entity recognition studies, the MetaMap (MMTx) program has been specially designed to take advantage of the features of the UMLS Metathesaurus and Specialist lexicon [33]. MMTx is available from <http://mmtx.nlm.nih.gov/>. Besides text mining, the Metathesaurus is used in a wide range of applications, including linking between different clinical or biomedical vocabularies, information retrieval and indexing, and biomedical language processing. The Metathesaurus is available from <http://umlsks.nlm.nih.gov/> (or on DVD) in relational database format. Users must complete the License Agreement for the Use of UMLS Metathesaurus. Java and XML application programming interfaces are available for the Metathesaurus.

3.5 Ontological Resources

Biomedical ontology aims to study the kinds of entities (i.e., substances, qualities, and processes) of biomedical significance. Unlike biomedical terminology, biomedical ontology is not primarily concerned with names, but with the principled definition of biological classes and their interrelations. In practice, since most terminologies have some degree of organization, and many ontologies also collect names for their entities, the distinction between ontological and terminological resources is somewhat arbitrary. See Chapter 4 for further discussion of this issue. Because they share many characteristics with ontologies, we will list under this rubric two broad resources—SNOMED CT and the UMLS Semantic Network. Other ontologies will be briefly discussed.

3.5.1 SNOMED CT

The Systematized Nomenclature of Medicine (SNOMED) Clinical Terms (SNOMED CT), developed by the College of American Pathologists, was formed by the convergence of SNOMED RT and Clinical Terms Version 3 (formerly known as the Read Codes). SNOMED CT is the most comprehensive biomedical terminology recently developed in native description logic formalism.¹ The version described here is dated January 31, 2004, and contains some 270,000 concepts, named by over 400,000 names. SNOMED CT consists of 18 independent hierarchies reflecting, in part, the organization of previous versions of SNOMED into “axes,” such as ‘Diseases’, ‘Drugs’, ‘Living organisms’, ‘Procedures’, and ‘Topography’. Each SNOMED CT concept is described by a variable number of elements. For example, the concept ‘Neurofibromatosis, type 2’ has a unique identifier (92503002), several names (*Bilateral acoustic neurofibromatosis*, *BANF - Bilateral acoustic neurofibromatosis*, *Neurofibromatosis, type 2*, and *Neurofibromatosis type 2*), and has multiple *is a* parents, including ‘Congenital anomaly of inner ear’, ‘Neoplasm of uncertain behavior of cranial nerve’, and ‘Acoustic neuroma’. ‘Neurofibromatosis, type 2’ also participates in a complex network of associative relations to other concepts. The relations (called roles), shown in Table 3.3, indicate that the lesions encountered in ‘Neurofibromatosis, type 2’ include neurofibromatosis of the vestibulocochlear nerve (Group 1) and neurilemoma of the vestibular nerve (Group 3). SNOMED CT is available as part of the UMLS (from <http://umlsks.nlm.nih.gov/>), at no charge for UMLS licensees in the United States. The structure of the UMLS Metathesaurus has been modified to accommodate the level of detail provided by ontological resources like SNOMED CT. Because SNOMED CT has become available through the UMLS only since 2004, the number of studies reporting its uses is still limited.

3.5.2 UMLS Semantic Network

The UMLS Semantic Network is one of three knowledge sources developed and distributed by the NLM, as part of the UMLS project. It was created in an effort to provide a semantic framework for the UMLS and its constituent vocabularies [34]. Unlike the Metathesaurus, the Semantic Network is a small structure, composed of 135 high-level categories called semantic types. It is organized in two single-inheritance hierarchies: one for ‘Entity’ and one for ‘Event’. In addition to *is a*, 53 kinds of relationships are defined in the Semantic Network, which are used to represent over 6,700 relations—hierarchical and associative—among semantic types. Semantic types from the Semantic Network are

1. Although not distributed in any description logic (DL) language, SNOMED CT has been developed in native DL formalism, thus ensuring the consistency of its hierarchical relations.

Table 3.3

Some of the Roles Present in the Definition of 'Neurofibromatosis, type 2'

Group	Role	Value
1	Associated morphology	'Neurofibromatosis'
	Finding site	'Skin structure'
	Finding site	'Vestibulocochlear nerve structure'
3	Associated morphology	'Neurilemoma'
	Finding site	'Vestibular nerve structure'

linked to Metathesaurus concepts by the categorization link established by the Metathesaurus editors. Each concept is categorized with at least one semantic type from the Semantic Network, independently of its hierarchical position in the source vocabularies. Fifteen collections of semantic types, called semantic groups, have been defined in order to partition Metathesaurus concepts into a smaller number of semantically consistent groups [35].

Semantic types for the Metathesaurus concepts listed in Table 3.1 are presented in Table 3.4, along with the corresponding semantic groups. For example, the concept 'Neurofibromatosis 2' is categorized as 'Neoplastic Process', a semantic type from the semantic group 'Disorders'. In addition to 'mutation,' Metathesaurus concepts categorized with 'Genetic Function' include 'alternative splicing', 'loss of heterozygosity', and 'ribonuclease activity'. Examples of

Table 3.4

Semantic Types and Semantic Groups for the Metathesaurus Concepts Listed in Table 3.1

ST Abbreviation	ST Name	Semantic Group
aapp	Amino Acid, Peptide, or Protein	Chemicals & Drugs
bacs	Biologically Active Substance	Chemicals & Drugs
bdsy	Body System	Anatomy
bpoc	Body Part, Organ, or Organ Component	Anatomy
celc	Cell Component	Anatomy
dsyn	Disease or Syndrome	Disorders
genf	Genetic Function	Physiology
gngm	Gene or Genome	Genes & Molecular Sequences
neop	Neoplastic Process	Disorders

relations among semantic types include ‘Body Part, Organ, or Organ Component’ *location of* ‘Neoplastic Process;’ ‘Pharmacologic Substance’ *treats* ‘Neoplastic Process;’ and ‘Neoplastic Process’ *manifestation of* ‘Genetic Function.’ A relationship between two semantic types indicates a *possible* link between the concepts categorized with these semantic types. In natural language processing and text mining applications, Semantic Network relations are typically used as supporting evidence for the candidate predicates (i.e., <concept₁, relationship, concept₂> structures) extracted from the text [36]. For example, in *schwannomas of cranial nerves*, after identifying the concepts ‘neurilemmoma’ (from *schwannoma*) as a ‘Neoplastic Process’ and *cranial nerves* as a ‘Body Part, Organ, or Organ Component,’ the preposition *of* can be interpreted as indicating the location of the neoplastic process to the body part. This candidate predicate is supported by the Semantic Network relation ‘Body Part, Organ, or Organ Component’ location of ‘Neoplastic Process.’ Many relation extraction systems rely on correspondences established between semantic relations and linguistic phenomena [37]. Semantic Network relations also can be exploited in conjunction with relations among concepts in the Metathesaurus [38]. The Semantic Network is distributed as part of the UMLS, and is available from <http://umlsks.nlm.nih.gov/>. Like the other UMLS knowledge sources, it can be queried through application programming interfaces for Java and XML.

3.5.3 Other Ontological Resources

In addition to SNOMED CT and the UMLS Semantic Network, several ontological resources can be used to support text mining. The Foundational Model of Anatomy (FMA) (<http://fma.biostr.washington.edu/>) is a large reference ontology of anatomy, developed at the University of Washington [39]. In addition to NLP applications [40], the FMA has been used in entity recognition tasks [41], as well as relation extraction tasks [42]. Ontologies, such as OpenGALEN (<http://www.opengalen.org/>), have been developed to support terminological services [43], and may be less useful for text mining applications. For example, unlike terminologies, OpenGALEN does not record lists of synonyms for biomedical entities. For more information about biomedical ontologies, we refer the interested reader to [44].

3.6 Issues Related to Entity Recognition

The biomedical domain has a long tradition of collecting and organizing terms, as well as building classifications, dating back to the seventeenth century. The dozens of terminological resources resulting from this effort now benefit entity

recognition tasks. Moreover, UMLS has contributed to make existing terminologies both easier to use, by providing a common format and distribution mechanism, and more useful, by identifying synonymy and other semantic relations across them. As part of this effort, the NLM also developed the lexical resources (lexicon and lexical programs) used to detect lexical similarity among biomedical terms, and, more generally, to process biomedical text. This is the reason why the UMLS is used in a large number of text mining systems in biomedicine.

The properties of biomedical terms have been studied. For example, [45, 46] found matches for 10% to 34% of the UMLS strings in MEDLINE/PubMed (depending on the matching criteria used), and [45] developed a model for identifying the UMLS terms useful in NLP applications. In the domain of molecular biology, researchers have investigated the lexical properties of the GO: 35% of GO terms have been found in the biomedical literature [31], and 66% of GO terms are composed of other GO terms [47]. A model of compositionality in GO has even been proposed [48]. These studies have confirmed the interest of using existing terminological resources in entity recognition tasks.

However, there are some remaining challenges in biomedical entity recognition, including the limited coverage of terminological resources and ambiguity in biomedical names.

3.6.1 Limited Coverage

First, some subdomains remain only partially covered by existing resources. One example is given by genes and proteins, and, more generally, chemical entities. Names for such entities have proved difficult to exhaustively compile in terminologies. Vocabularies extracted from specialized databases may complement traditional terminologies. Moreover, while variant formation has been studied and effectively modeled for clinical terms [49], normalization techniques for the less regular names of entities employed in genomics have been only recently researched [50]. For these reasons, entity recognition techniques in this subdomain often include machine learning approaches, rather than the rule-based approach traditionally employed in biomedical NLP. Many gene name identification systems have been developed in the last 5 years [51–54]. Entity recognition systems in molecular biology texts may include algorithms rather than (or in addition to) static resources [24]. However, the product of some of these algorithms is made available to the research community by their authors. For example, [55] shares the lexicon of over 1 million gene and protein names extracted from the biomedical literature. Coverage issues have been explored in clinical terminologies as well [56], and techniques have been developed to extend the coverage of terminologies to specialized subdomains [57], or from specific corpora [58]. More generally, relation

extraction also may benefit from term extraction techniques resulting from research in terminology [59].

3.6.2 Ambiguity

The second issue is the ambiguity of many names in biology. This phenomenon is common in natural language, but poses specific challenges to biomedical entity recognition. Polysemy (several meanings for the same name) is illustrated by *NF2*, which simultaneously names the gene, the protein it produces, and the disease resulting from its mutation. While polysemy usually does not pose problems for domain experts, it makes it difficult for entity recognition systems to select the appropriate meaning. The ambiguity resulting from polysemous gene names has been quantified by Chen et al. [60]. These authors found modest ambiguities with general English words (0.57%) and medical terms (1.01%), but high ambiguity across species (14.20%). Ambiguity across species may be difficult to resolve; for example, when only capitalization conventions differentiate between gene names in various model organisms (e.g., *NF2* in *Homo sapiens* versus *Nf2* in *Mus musculus*). Various disambiguation strategies have been applied to biomedical language processing [61, 62]. However, further research is needed to develop strategies adapted to the specificity of molecular biology (e.g., ambiguity across species). Moreover, the limited availability of annotated resources, such as the GENIA corpus [63], hinders the development of unsupervised disambiguation techniques.

3.7 Issues Related to Relation Extraction

We round off our discussion of resources with a brief look at issues concerning relation extraction, to further emphasize the core enabling role of resources with respect to text mining.

3.7.1 Terminological Versus Ontological Relations

Not only do terminologies contain a large number of names for biomedical entities useful for entity recognition tasks, but they also represent a similarly considerable number of relations. For example, over 16 million relations are recorded in the UMLS Metathesaurus. While not all of them represent well-defined predicates or assertions, as would be expected from ontologies, these relations are essentially beneficial to applications such as relation extraction, especially when used in combination with lexico-syntactic clues and additional ontological relations.

The relations found in the most recent terminologies—often developed using knowledge representation techniques such as description logics—are generally better specified and principled, and therefore more directly useful for relation extraction. However, a careful inspection of these and other ontological resources through the prism of formal ontology reveals some limitations, especially in terms of consistency [15, 64, 65]. Applying formal ontological principles to biomedical ontologies results in a clarification of the relations [66], which, in turn, is expected to result in more consistent ontologies and more accurate inferences.

Recent experiments in reengineering terminologies have shown both the benefit and the cost (in terms of human resources) of such efforts [67, 68]. However, improving ontologies is likely to benefit relation extraction, as the candidate assertions extracted from text must be checked, not necessarily against relations explicitly represented in ontologies, but most often against inferred relations.

3.7.2 Interactions Between Text Mining and Terminological Resources

This chapter deliberately looks at ontologies and other resources as enabling resources for text mining, and relation extraction in particular. Conversely, it is worth mentioning that the relations extracted from text corpora and other knowledge sources (e.g., annotation databases) can help identify additional ontological relations. For example, lexico-syntactic patterns have been used to extract hypernymy relations from text corpora [69], and statistical methods have helped identify associative relations among GO terms [70]. In other words, the relations between text mining techniques and terminological resources are not unilateral. There is a virtual cycle, in which applications and resources benefit from one another. Studying this symbiotic relation is beyond the scope of this chapter. Various existing resources can be combined in order to create new resources. For example, semantic lexicons have been derived from lexicons, terminologies, and text corpora [71, 72].

3.8 Conclusion

This chapter presents the various kinds of enabling resources used in biomedical text mining applications. Lexicons support basic natural language processing tasks, such as parsing. Along with terminologies, lexicons also provide lists of names (including variants) for biological entities, thus supporting entity recognition tasks. Finally, the relations represented in ontologies and terminologies often serve as a reference for relation extraction algorithms.

Because it integrates these three kinds of resources, the UMLS plays a central role in biomedical text mining. Its three components (Specialist lexicon, Metathesaurus, and Semantic Network) are illustrated in this chapter, showing their use in entity recognition and relation extraction tasks. The role of other resources, either more specialized or more general, is also discussed.

Despite the existence of these resources, there remain many challenges to entity recognition and relation extraction in biology. Existing biomedical lexicons and terminologies fail to provide adequate coverage of specialized subdomains (e.g., genes and proteins for the various model organisms). Approaches to normalizing the names of genomic entities and to resolving the ambiguity introduced by some of them need to be further researched. Finally, the development of large, consistent, principled sources of biomedical knowledge—namely ontologies—will benefit not only text mining applications, but, more generally, the wide range of tasks relying upon biomedical knowledge (e.g., database interoperability, decision support, and so forth).

Acknowledgments

The author would like to thank Tom Rindflesch, Pierre Zweigenbaum, and Karin Verspoor for useful comments on a previous version of this manuscript.

References

- [1] de Bruijn, B., and J. Martin, "Getting to the (C)Ore of Knowledge: Mining Biomedical Literature," *Int. J. Med. Inform.*, Vol. 67, No. 1–3, 2002, pp. 7–18.
- [2] Shatkay, H., and R. Feldman, "Mining the Biomedical Literature in the Genomic Era: An Overview," *J. Comput. Biol.*, Vol. 10, No. 6, 2003, pp. 821–855.
- [3] Yandell, M. D., and W. H. Majoros, "Genomics and Natural Language Processing," *Nat. Rev. Genet.*, Vol. 3, No. 8, 2002, pp. 601–610.
- [4] Hahn, U., M. Romacker, and S. Schulz, "Medsyndikate—A Natural Language System for the Extraction of Medical Information from Findings Reports," *Int. J. Med. Inform.*, Vol. 67, No. 1–3, 2002, pp. 63–74.
- [5] Corney, D. P., et al., "BioRAT: Extracting Biological Information from Full-Length Papers," *Bioinformatics*, Vol. 20, No. 17, 2004, pp. 3206–3213.
- [6] Leroy, G., H. Chen, and J. D. Martinez, "A Shallow Parser Based on Closed-Class Words to Capture Relations in Biomedical Text," *J. Biomed. Inform.*, Vol. 36, No. 3, 2003, pp. 145–158.
- [7] Humphreys, K., G. Demetriou, and R. Gaizauskas, "Two Applications of Information Extraction to Biological Science Journal Articles: Enzyme Interactions and Protein Structures," *Proc. Pacific Symp. on Biocomputing*, 2000, pp. 505–516.

- [8] Friedman, C., et al., "A General Natural-Language Text Processor for Clinical Radiology," *J. Am. Med. Inform. Assoc.*, Vol. 1, No. 2, 1994, pp. 161–74.
- [9] Friedman, C., et al., "GENIES: A Natural-Language Processing System for the Extraction of Molecular Pathways from Journal Articles," *Bioinformatics*, Vol. 17, Suppl. 1, 2001, pp. S74–82.
- [10] Rindflesch, T. C., M. Fiszman, and B. Libbus, "Semantic Interpretation for the Biomedical Literature," *Medical Informatics: Advances in Knowledge Management and Data Mining in Biomedicine*, H. Chen, et al., (eds.), Berlin: Springer-Verlag, 2005, pp. 399–422.
- [11] Grabar, N., et al., "Matching Controlled Vocabulary Words," *Stud. Health Technol. Inform.*, Vol. 95, 2003, pp. 445–450.
- [12] Bodenreider, O., and C. A. Bean, "Relationships among Knowledge Structures: Vocabulary Integration Within a Subject Domain," *Relationships in the Organization of Knowledge*, C. A. Bean, and R. Green, (eds.), Amsterdam: Kluwer, 2001, pp. 81–98.
- [13] Nenadic, G., I. Spasic, and S. Ananiadou, "Terminology-Driven Mining of Biomedical Literature," *Bioinformatics*, Vol. 19, No. 8, 2003, pp. 938–943.
- [14] Smith, B., et al., "Relations in Biomedical Ontologies," *Genome Biol.*, Vol. 6, No. 5, 2005, p. R46.
- [15] Smith, B., J. Williams, and S. Schulze-Kremer, "The Ontology of the Gene Ontology," *Proc. AMIA Annual Symp.*, 2003, pp. 609–613.
- [16] Bodenreider, O., "The Unified Medical Language System (UMLS): Integrating Biomedical Terminology," *Nucleic Acids Res.*, Vol. 32, Database issue, 2004, pp. D267–D270.
- [17] Baser, M. E., D. G. Evans, and D. H. Gutmann, "Neurofibromatosis 2," *Curr. Opin. Neurol.*, Vol. 16, No. 1, 2003, pp. 27–33.
- [18] Uppal, S., and A. P. Coatesworth, "Neurofibromatosis Type 2," *Int. J. Clin. Pract.*, Vol. 57, No. 8, 2003, pp. 698–703.
- [19] Fellbaum, C., *WordNet: An Electronic Lexical Database, Language, Speech, and Communication*. Cambridge, MA: MIT Press, 1998.
- [20] Bodenreider, O., A. Burgun, and J. A. Mitchell, "Evaluation of WordNet as a Source of Lay Knowledge for Molecular Biology and Genetic Diseases: A Feasibility Study," *Stud. Health Technol. Inform.*, Vol. 95, 2003, pp. 379–384.
- [21] Burgun, A., and O. Bodenreider, "Comparing Terms, Concepts and Semantic Classes in WordNet and the Unified Medical Language System," *Proc. NAACL2001 Workshop, WordNet and Other Lexical Resources: Applications, Extensions and Customizations*, 2001, pp. 77–82.
- [22] Leroy, G., and H. Chen, "Meeting Medical Terminology Needs—The Ontology-Enhanced Medical Concept Mapper," *IEEE Trans. on Inf. Technol. Biomed.*, Vol. 5, No. 4, 2001, pp. 261–270.
- [23] Browne, A. C., et al., "UMLS Language and Vocabulary Tools," *Proc. AMIA Annual Symp.*, 2003, p. 798.
- [24] Krauthammer, M., and G. Nenadic, "Term Identification in the Biomedical Literature," *J. Biomed. Inform.*, Vol. 37, No. 6, 2004, pp. 512–526.

- [25] Wain, H. M., et al., "Genew: The Human Gene Nomenclature Database, 2004 Updates," *Nucleic Acids Res.*, Vol. 32, Database issue, 2004, pp. D255–D257.
- [26] Pustejovsky, J., et al., "Automatic Extraction of Acronym-Meaning Pairs from MEDLINE Databases," *Medinfo*, Vol. 10, No. 1, 2001, pp. 371–375.
- [27] Schwartz, A. S., and M. A. Hearst, "A Simple Algorithm for Identifying Abbreviation Definitions in Biomedical Text," *Proc. Pacific Symp. on Biocomputing*, 2003, pp. 451–462.
- [28] Wren, J. D., et al., "Biomedical Term Mapping Databases," *Nucleic Acids Res.*, Vol. 33, Database Issue, 2005, pp. D289–D293.
- [29] Chute, C. G., "Clinical Classification and Terminology: Some History and Current Observations," *J. Am. Med. Inform. Assoc.*, Vol. 7, No. 3, 2000, pp. 298–303.
- [30] Ashburner, M., et al., "Gene Ontology: Tool for the Unification of Biology: The Gene Ontology Consortium," *Nat. Genet.*, Vol. 25, No. 1, 2000, pp. 25–29.
- [31] McCray, A. T., A. C. Browne, and O. Bodenreider, "The Lexical Properties of the Gene Ontology," *Proc. AMIA Annual Symp.*, 2002, pp. 504–508.
- [32] Nelson, S. J., D. Johnston, and B. L. Humphreys, "Relationships in Medical Subject Headings," *Relationships in the Organization of Knowledge*, C. A. Bean and R. Green, (eds.), Amsterdam: Kluwer, 2001, pp. 171–184.
- [33] Aronson, A. R., "Effective Mapping of Biomedical Text to the UMLS Metathesaurus: The Metamap Program," *Proc. AMIA Annual Symp.*, 2001, pp. 17–21.
- [34] McCray, A. T., "An Upper-Level Ontology for the Biomedical Domain," *Comp. Funct. Genom.*, No. 4, 2003, pp. 80–84.
- [35] Bodenreider, O., and A. T. McCray, "Exploring Semantic Groups Through Visual Approaches," *J. Biomed. Inform.*, Vol. 36, No. 6, 2003, pp. 414–432.
- [36] Rindflesch, T. C., and M. Fisman, "The Interaction of Domain Knowledge and Linguistic Structure in Natural Language Processing: Interpreting Hypernymic Propositions in Biomedical Text," *J. Biomed. Inform.*, Vol. 36, No. 6, 2003, pp. 462–477.
- [37] Libbus, B., et al., "Using Natural Language Processing, LocusLink, and the Gene Ontology to Compare OMIM to MEDLINE," *Proc. HLT-NAACL Workshop on Linking the Biological Literature, Ontologies and Databases: Tools for Users*, 2004, pp. 69–76.
- [38] McCray, A. T., and O. Bodenreider, "A Conceptual Framework for the Biomedical Domain," *The Semantics of Relationships: An Interdisciplinary Perspective*, R. Green, C. A. Bean, and S. H. Myaeng, (eds.), Boston, MA: Kluwer Academic Publishers, 2002, pp. 181–198.
- [39] Rosse, C., and J. L. Mejino, Jr., "A Reference Ontology for Biomedical Informatics: The Foundational Model of Anatomy," *J. Biomed. Inform.*, Vol. 36, No. 6, 2003, pp. 478–500.
- [40] Distelhorst, G., et al., "A Prototype Natural Language Interface to a Large Complex Knowledge Base, the Foundational Model of Anatomy," *Proc. AMIA Annual Symp.*, 2003, pp. 200–204.

- [41] Sneiderman, C. A., T. C. Rindfleisch, and C. A. Bean, "Identification of Anatomical Terminology in Medical Text," *Proc. AMIA Annual Symp.*, 1998, pp. 428–432.
- [42] Bean, C. A., T. C. Rindfleisch, and C. A. Sneiderman, "Automatic Semantic Interpretation of Anatomic Spatial Relationships in Clinical Text," *Proc. AMIA Annual Symp.*, 1998, pp. 897–901.
- [43] Nowlan, W. A., et al., "From Terminology to Terminology Services," *Proc. Annual Symp. on Computer Applications in Medical Care*, 1994, pp. 150–154.
- [44] Bodenreider, O., and A. Burgun, "Biomedical Ontologies," *Medical Informatics: Advances in Knowledge Management and Data Mining in Biomedicine*, H. Chen, et al., (eds.), Berlin: Springer-Verlag, 2005, pp. 211–236.
- [45] McCray, A. T., et al., "Evaluating UMLS Strings for Natural Language Processing," *Proc. AMIA Annual Symp.*, 2001, pp. 448–452.
- [46] Srinivasan, S., et al., "Finding UMLS Metathesaurus Concepts in MEDLINE," *Proc. AMIA Annual Symp.*, 2002, pp. 727–731.
- [47] Ogren, P. V., et al., "The Compositional Structure of Gene Ontology Terms," *Proc. Pacific Symp. on Biocomputing*, 2004, pp. 214–525.
- [48] Mungall, C., "Obol: Integrating Language and Meaning in Bio-Ontologies," *Comparative and Functional Genomics*, Vol. 5, No. 7, 2004, pp. 509–520.
- [49] McCray, A. T., S. Srinivasan, and A. C. Browne, "Lexical Methods for Managing Variation in Biomedical Terminologies," *Proc. Annual Symp. on Computer Applications in Medical Care*, 1994, pp. 235–239.
- [50] Morgan, A. A., et al., "Gene Name Identification and Normalization Using a Model Organism Database," *J. Biomed. Inform.*, Vol. 37, No. 6, 2004, pp. 396–410.
- [51] Collier, N., and K. Takeuchi, "Comparison of Character-Level and Part of Speech Features for Name Recognition in Biomedical Texts," *J. Biomed. Inform.*, Vol. 37, No. 6, 2004, pp. 423–435.
- [52] Koike, A., Y. Niwa, and T. Takagi, "Automatic Extraction of Gene/Protein Biological Functions from Biomedical Text," *Bioinformatics*, Vol. 21, No. 7, 2005, pp. 1227–1236.
- [53] Proux, D., et al., "Detecting Gene Symbols and Names in Biological Texts: A First Step Toward Pertinent Information Extraction," *Genome Inform. Ser.*, Vol. 9, 1998, pp. 72–80.
- [54] Yu, H., et al., "Automatic Extraction of Gene and Protein Synonyms from MEDLINE and Journal Articles," *Proc. AMIA Annual Symp.*, 2002, pp. 919–923.
- [55] Tanabe, L., and W. J. Wilbur, "Generation of a Large Gene/Protein Lexicon by Morphological Pattern Analysis," *J. Bioinform. Comput. Biol.*, Vol. 1, No. 4, 2004, pp. 611–26.
- [56] Chute, C. G., et al., "The Content Coverage of Clinical Classifications: For the Computer-Based Patient Record Institute's Work Group on Codes & Structures," *J. Am. Med. Inform. Assoc.*, Vol. 3, No. 3, 1996, pp. 224–233.
- [57] Harris, M. R., et al., "A Term Extraction Tool for Expanding Content in the Domain of Functioning, Disability, and Health: Proof of Concept," *J. Biomed. Inform.*, Vol. 36, No. 4–5, 2003, pp. 250–259.

- [58] Bodenreider, O., T. C. Rindfleisch, and A. Burgun, "Unsupervised, Corpus-Based Method for Extending a Biomedical Terminology," *Proc. ACL'2002 Workshop Natural Language Processing in the Biomedical Domain*, 2002, pp. 53–60.
- [59] Jacquemin, C., *Spotting and Discovering Terms Through Natural Language Processing*, Cambridge, MA: MIT Press, 2001.
- [60] Chen, L., H. Liu, and C. Friedman, "Gene Name Ambiguity of Eukaryotic Nomenclatures," *Bioinformatics*, Vol. 21, No. 2, 2005, pp. 248–256.
- [61] Liu, H., Y. A. Lussier, and C. Friedman, "Disambiguating Ambiguous Biomedical Terms in Biomedical Narrative Text: An Unsupervised Method," *J. Biomed. Inform.*, Vol. 34, No. 4, 2001, pp. 249–261.
- [62] Liu, H., V. Teller, and C. Friedman, "A Multi-Aspect Comparison Study of Supervised Word Sense Disambiguation," *J. Am. Med. Inform. Assoc.*, Vol. 11, No. 4, 2004, pp. 320–331.
- [63] Kim, J. D., et al., "Genia Corpus—Semantically Annotated Corpus for Bio-Textmining," *Bioinformatics*, Vol. 19, Suppl. 1, 2003, pp. i180–i182.
- [64] Ceusters, W., et al., "Ontology-Based Error Detection in SNOMED-Ct(R)," *Medinfo*, Vol. 2004, 2004, pp. 482–486.
- [65] Smith, B., A. Kumar, and S. Schulze-Kremer, "Revising the UMLS Semantic Network," *Medinfo*, Vol. 2004, No. CD, 2004, p. 1700.
- [66] Smith, B., and C. Rosse, "The Role of Foundational Relations in the Alignment of Biomedical Ontologies," *Medinfo*, Vol. 2004, pp. 444–448.
- [67] Hahn, U., and S. Schulz, "Towards a Broad-Coverage Biomedical Ontology Based on Description Logics," *Proc. Pacific Symp. on Biocomputing*, 2003, pp. 577–588.
- [68] Wroe, C. J., et al., "A Methodology to Migrate the Gene Ontology to a Description Logic Environment Using DAML+OIL," *Proc. Pacific Symp. on Biocomputing*, 2003, pp. 624–635.
- [69] Hearst, M. A., "Automatic Acquisition of Hyponyms from Large Text Corpora," *Proc. 15th Int. Conf. on Computational Linguistics (COLING-92)*, 1992, pp. 539–545.
- [70] Bodenreider, O., M. Aubry, and A. Burgun, "Non-Lexical Approaches to Identifying Associative Relations in the Gene Ontology," *Proc. Pacific Symp. on Biocomputing*, 2005, pp. 104–115.
- [71] Johnson, S. B., "A Semantic Lexicon for Medical Language Processing," *J. Am. Med. Inform. Assoc.*, Vol. 6, No. 3, 1999, pp. 205–218.
- [72] Verspoor, K., "Towards a Semantic Lexicon for Biological Language Processing," *Comparative and Functional Genomics*, Vol. 6, No. 1–2, 2005, pp. 61–66.