

MedTag: A Collection of Biomedical Annotations

L.H. Smith[†], L. Tanabe[†], T. Rindflesch[‡], W.J. Wilbur[†]

[†]National Center for Biotechnology Information

[‡]Lister Hill National Center for Biomedical Communications

NLM, NIH, 8600 Rockville Pike, Bethesda, MD 20894

{lsmith,tanabe,wilbur}@ncbi.nlm.nih.gov

rindflesch@nlm.nih.gov

Abstract

We present a database of annotated biomedical text corpora merged into a portable data structure with uniform conventions. MedTag combines three corpora, MedPost, ABGene and GENETAG, within a common relational database data model. The GENETAG corpus has been modified to reflect new definitions of genes and proteins. The MedPost corpus has been updated to include 1,000 additional sentences from the clinical medicine domain. All data have been updated with original MEDLINE text excerpts, PubMed identifiers, and tokenization independence to facilitate data accuracy, consistency and usability.

The data are available in flat files along with software to facilitate loading the data into a relational SQL database from <ftp://ftp.ncbi.nlm.nih.gov/pub/lsmith/MedTag/medtag.tar.gz>.

1 Introduction

Annotated text corpora are used in modern computational linguistics research and development to fine-tune computer algorithms for analyzing and classifying texts and textual components. Two important factors for useful text corpora are 1) accuracy and consistency of the annotations, and 2) usability of the data. We have recently updated the text corpora we use in our research with respect to these criteria.

Three different corpora were combined. The ABGene corpus consists of over 4 000 sentences annotated with gene and protein named entities. It was originally used to train the ABGene tagger to recognize gene/protein names in MEDLINE records, and recall and precision rates in the lower 70 percentile range were achieved (Tanabe and Wilbur, 2002). The MedPost corpus consists of 6 700 sentences, and is annotated with parts of speech, and gerund arguments. The MedPost tagger was trained on 3 700 of these sentences and achieved an accuracy of 97.4% on the remaining sentences (Smith et. al., 2004). The GENETAG corpus for gene/protein named entity identification, consists of 20 000 sentences and was used in the BioCreative 2004 Workshop (Yeh et. al., 2005; Tanabe et. al., 2005) (only 15 000 sentences are currently released, the remaining 5 000 are being retained for possible use in a future workshop). Training on a portion of the data, the top performing systems achieved recall and precision rates in the lower 80 percentile range. Because of the scarcity of good annotated data in the realm of biomedicine, and because good performance has been obtained using this data, we feel there is utility in presenting it to a wider audience.

All of the MedTag corpora are based on MEDLINE abstracts. However, they were queried at different times, and used different (but similar) algorithms to perform tokenization and sentence segmentation. The original annotations were assigned to tokens, or sequences of tokens, and extensively reviewed by the authors at different times for the different research projects.

The main goals in combining and updating these

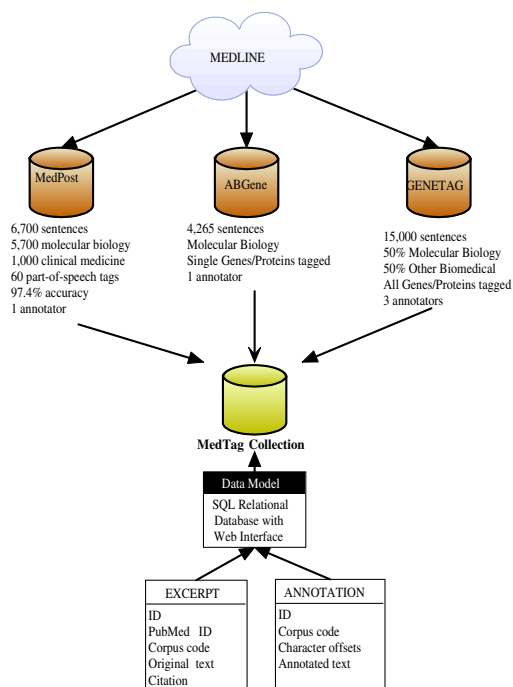


Figure 1: Component corpora, common data model and main record types of the MedTag collection.

corpora into a single corpus were to

1. update the text for all corpora to that currently found in MEDLINE, storing a correct citation and the original, untokenized text for each excerpt
2. eliminate tokenization dependence
3. put all text and annotations into a common database format
4. provide programs to convert from the new corpus format to the data formats used in previous research

2 Merging the Corpora

We describe what was done to merge the original corpora, locating original sources and modifying the text where needed. An overview is given in Figure 1. Some basic statistics are given in Table 1.

2.1 Identifying Source Data

The original data of the three corpora were assembled and the text was used to search MEDLINE to

Corpus	sentences	tokens	most frequent tag
GENETAG-05	15,000	418,246	insulin_GENE(112)
MedPost	6,700	181,626	the_DD(8,507)
ABGene	4,265	123,208	cyclin_GENE(165)

MedPost

Adj	Adv	Aux	Noun	Punct	Verb
14,648	4,553	56,262	60,732	21,806	23,625

GENETAG-05

GENE	ALTGENE
24,562	19,216

ABGene

GENE	ALTGENE
8,185	0

Table 1: MedTag Corpora. GENE = gene and protein names, ALTGENE = acceptable alternatives for gene and protein names. MedPost tagset contains 60 parts of speech which have been binned here for brevity.

find the closest match. An exact or near exact match was found for all but a few excerpts. For only a few excerpts, the MEDLINE record from which the excerpt was originally taken had been removed or modified and an alternative sentence was selected. Thus, each excerpt in the database is taken from a MEDLINE record as it existed at one time in 2004. In order to preserve the reference for future work, the PubMed ID and citation data were also retrieved and stored with each excerpt. Each excerpt in the current database roughly corresponds to a sentence, although the procedure that extracted the sentence is not specified.

2.2 Eliminating Tokenization Dependence

In the original ABGene and GENETAG corpora, the gene and protein phrases were specified by the tokens contained in the phrase, and this introduced a dependence on the tokenization algorithm. This created problems for researchers who wished to use a different tokenization. To overcome this dependence, we developed an alternative way of specifying

ing phrases. Given the original text of an excerpt, the number of non-whitespace characters to the start of the phrase does not depend on the tokenization. Therefore, all annotations now refer to the first and last character of the phrase that is annotated. For example the protein *serum LH* in the excerpt

There was no correlation between *serum LH* and chronological or bone age in this age group, which suggests that the correlation found is not due to age-related parallel phenomena.

is specified as characters 28 to 34 (the first character is 0).

2.3 Data Model

There are two main record types in the database, EXCERPT and ANNOTATION. Each EXCERPT record stores an identifier and the original corpus code (abgene, medpost, and genetag) as well as sub-corpus codes that were defined in the original corpora. The original text, as it was obtained from MEDLINE, is also stored, and a human readable citation to the article containing the reference.

Each ANNOTATION record contains a reference to the excerpt (by identifier and corpus), the character offset of the first and last characters of the phrase being annotated (only non-whitespace characters are counted, starting with 0), and the corresponding annotation. The annotated text is stored for convenience, though it can be obtained from the corresponding excerpt record by counting non-whitespace characters.

The data is provided as an ASCII file in a standard format that can be read and loaded into a relational database. Each record in the file begins with a line of the form `>>table_name` where *table_name* is the name of the table for that record. Following the table name is a series of lines with the form *field: value* where *field* is the name of the field and *value* is the value stored in that field.

Scripts are provided for loading the data into a relational database, such as *mysql* or *ORACLE*. SQL queries can then be applied to retrieve excerpts and annotations satisfying any desired condition. For example, here is an SQL query to retrieve excerpts from the MedPost corpus containing the token *p53* and *signaling* or *signalling*

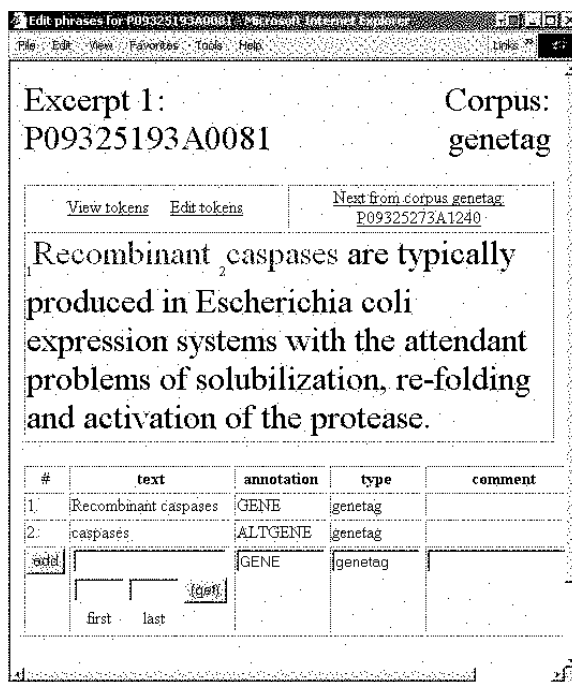


Figure 2: A screen capture of the annotator's interface and the GENETAG-05 annotations for a sentence.

```
select text from excerpt
where text like '%p53%'
and text rlike 'signa[l]*ing';
```

2.4 Web Interface

A web-based corpus editor was used to enter and review annotations. The code is being made available, as is, and requires that the data are loaded into a mysql database that can be accessed by a web server. The interface supports two annotation types: MedPost tags and arbitrary phrase annotations. MedPost tags are selectable from a pull-down menu of pre-programmed likely tags. For entering phrase annotations, the user highlights the desired phrase, and pressing the enter key computes and saves the first and last character offsets. The user can then enter the annotation code and an optional comment before saving it in the database. A screen dump of the phrase annotations for a sentence in the genetag corpus is shown in figure 2.

The data from the database was dumped to the flat file format for this release. We have also included some files to accommodate previous users of the corpora. A perl program, *alt_eval.perl* is in-

cluded that replaces the GENETAG evaluation program using non-whitespace character numbers instead of token numbers. Copies of the ABGene and MedPost corpora, in the original formats, are also included.

3 Updates of Component Corpora

3.1 MedPost Update

The MedPost corpus (Smith et. al., 2004) originally contained 5 700 tokenized sentences. An additional 1 000 annotated sentences have been added for this release. Each sentence in the MedPost corpus is fully tokenized, that is, divided into non-overlapping annotated portions, and each token is annotated with one of 60 part of speech tags (see Table 1). Minor corrections to the annotations have been made since the original release.

Since most of the original corpus, and all of the sentences used for training the MedPost tagger, were in the area of molecular biology, we added an additional 1 000 sentences selected from random MEDLINE abstracts on the subject of clinical medicine. As a preliminary result, the trained MedPost tagger achieves approximately 96.9% accuracy, which is comparable to the 97.4% accuracy achieved on the subset of 1 000 sentences selected randomly from all of MEDLINE. An example of a sentence from the clinical medicine collection is

Evidence_{NN} is_{VBZ} now_{RR} available_{JJ}
to_{TO} show_{VVI} a_{DD} beneficial_{JJ} effect_{NN}
of_{II} bezafibrate_{NN} on_{II} retarding_{VVGN}
atherosclerotic_{JJ} processes_{NNS} and_{CC} in_{II}
reducing_{VVGN} risk_{NN} of_{II} coronary_{JJ} heart_{NN}
disease_{NN} .

In addition to the token-level annotations, all of the gerunds in the MedPost corpus (these are tagged *VVGN*) were also examined and it was noted whether the gerund had an explicit subject, direct object, or adjective complement. This annotation is stored with an annotation of type *gerund*. To illustrate, the two gerunds in the previous example, *retarding* and *reducing* both have direct objects (*retarding processes* and *reducing risk*), and the gerund tag is entered as “o”. The gerund annotations have been used to improve a noun phrase bracketer able to recognize gerundive phrases.

3.2 GENETAG Update

GENETAG is a corpus of MEDLINE sentences that have been annotated with gene and protein names. The closest related work is the GENIA corpus (Kim et. al., 2003). GENIA provides detailed coverage of a large number of semantic entities related to a specific subset of human molecular biology, whereas GENETAG provides gene and protein name annotations only, for a wide range of organisms and biomedical contexts (molecular biology, genetics, biochemistry, clinical medicine, etc.)

We are including a new version of GENETAG, GENETAG-05, as part of the MedTag system. GENETAG-05 differs from GENETAG in four ways: 1) the definition of a gene/protein entity has been modified, 2) significant annotation errors in GENETAG have been corrected, 3) the concept of a non-specific entity has been refined, and 4) character-based indices have been introduced to reduce tokenization problems. We believe that these changes result in a more accurate and robust corpus.

GENETAG-05 maintains a wide definition of a gene/protein entity including genes, proteins, domains, sites, sequences, and elements, but excluding plasmids and vectors. The specificity constraint requires that a gene/protein name must be included in the tagged entity. This constraint has been applied more consistently in GENETAG-05. Additionally, plain sequences like *ATTGGCCTT-TAAC* are no longer tagged, embedded names are tagged (*ras*-mediated), and significantly more terms have been judged to violate the specificity constraint (*growth factor*, *proteases*, *protein kinase*, *ribonuclease*, *snoRNA*, *rRNA*, *tissue factor*, *tumor antigen*, *complement*, *hormone receptors*, *nuclear factors*, etc.).

The original GENETAG corpus contains some entities that were erroneously tagged as gene/proteins. Many of these errors have been corrected in the updated corpus. Examples include *camp-responsive elements*, *mu element*, *VDRE*, *melanin*, *dentin*, *myelin*, *auxin*, *BARBIE box*, *carotenoids*, and *cellulose*. Error analysis resulted in the updated annotation conventions given in Table 1.

Enzymes are a special class of proteins that catalyze biochemical reactions. Enzyme names have varying degrees of specificity, so the line drawn for

tagging purposes is based on online resources¹ as well as background knowledge. In general, tagged enzymes refer to more specific entities than untagged enzymes (*tyrosine kinase* vs. *protein kinase*, *ATPase* vs. *protease*). Enzymes that can refer to either DNA or RNA are tagged if the reference is specified (*DNA endonuclease* vs. *endonuclease*). Enzymes that do not require DNA/RNA distinction are tagged (*lipase* vs. *ligase*, *cyclooxygenase* vs. *methylase*). Non-specific enzymes are tagged if they clearly refer to a gene or protein, as in (1).

- 1) The structural gene for *hydrogenase* encodes a protein product of molecular mass 45820 Da.

Semantic constraints in GENETAG-05 are the same as those for GENETAG. To illustrate, the name in (2) requires *rabies* because *RIG* implies that the gene mentioned in this sentence refers to the *rabies immunoglobulin*, and not just any *immunoglobulin*. In (3), the word *receptor* is necessary to differentiate *IGG receptor* from *IGG*, a crucial biological distinction. In (4), the number *1* is needed to accurately describe a specific type of *tumor necrosis factor*, although *tumor necrosis factor* alone might be adequate in a different context.

- 2) rabies immunoglobulin (RIG)
- 3) IGG receptor
- 4) Tumor necrosis factor 1

Application of the semantic constraint can result in apparent inconsistencies in the corpus (*immunoglobulin* is sufficient on its own in some sentences in the corpus, but is insufficient in (2)). However, we believe it is important that the tagged entity retain its true meaning in the *sentence context*.

4 Recommended Uses

We have found the component corpora of MedTag to be useful for the following functions:

- 1) Training and evaluating part-of-speech taggers
- 2) Training and evaluating gene/protein named entity taggers

¹<http://cancerweb.nci.ac.uk/omd/copyleft.html>
<http://www.onelook.com/>

- 3) Developing and evaluating a noun phrase bracketer for PubMed phrase indexing
- 4) Statistical analysis of grammatical usage in medical text
- 5) Feature generation for machine learning

The MedPost tagger was recently ported to Java and is currently being employed in MetaMap, a program that maps natural language text into the UMLS (Aronson,A.R., 2001).

5 Conclusion

We have merged three biomedical corpora into a collection of annotations called MedTag. MedTag uses a common relational database format along with a web interface to facilitate annotation consistency. We have identified the MEDLINE excerpts for each sentence and eliminated tokenization dependence, increasing the usability of the data. In GENETAG-05, we have clarified many grey areas for annotation, providing better guidelines for tagging these cases. For users of previous versions of the component corpora, we have included programs to convert from the new standardized format to the formats used in the older versions.

References

- Aronson, A. R. 2001. Effective mapping of biomedical text to the UMLS Metathesaurus: the MetaMap program. Proc. AMIA Symp., 1721.
- Kim, J.-D., Ohta, T., Tateisi, Y. and Tsujii, J. 2003. GENIA corpus: a semantically annotated corpus for biotextmining. *Bioinformatics*, 19: 180 - 182.
- Tanabe, L and Wilbur, WJ. 2002. Tagging gene and protein names in biomedical text. *Bioinformatics*, 18, 1124-1132.
- Tanabe L, Xie N, Thom, LH, Matten W, Wilbur, WJ: GENETAG: a tagged gene corpus for gene/protein named entity recognition. *BMC Bioinformatics* 2005.
- Smith, L, Rindfleisch, T, and Wilbur, WJ. 2004. MedPost: a part of speech tagger for biomedical text. *Bioinformatics*, 20(13) 2320-2321.
- Yeh A, Hirschman L, Morgan A, Colosimo M: BioCre-AtIvE task 1A: gene mention finding evaluation. *BMC Bioinformatics* 2005.

Entity Type	Problem	GENETAG-05 Convention	Positive Examples	Negative Examples
Protein Families	Some are named after structural motifs.	Do not tag structures alone, but tag structurally related gene and protein families.	<i>Zinc finger protein, bZIP transcription factor, homeobox gene, TATA binding protein</i>	<i>Zinc finger, helix-turn-helix motif, leucine zipper, homeobox, TATA box</i>
Domains	Name can refer to 1) the amino acid content of a sequence (<i>PEST</i>), 2) the protein that binds the sequence (<i>TFIIIA DNA binding domain</i>), 3) a homologous gene (<i>SH2 - Src homology domain 2</i>), 4) the first proteins in which the domain was discovered (<i>LIM, PDZ</i>), or 5) structural entities (<i>POZ, zinc finger domain</i>).	Tag only if the domain refers to a gene or protein. Immuno-globulin regions are tagged. (<i>VH</i> refers to the <i>Immuno-globulin heavy chain V region</i>).	<i>BTB domain, LIM domain, HECT domain, VH domain, SH2 domain, TFIIIA DNA binding domain, Krüppel-associated box (KRAB) domains, NF-IL6 beta leucine zipper domain</i>	<i>PEST domain, SR domain, zinc finger domain, b-Zip domain, POZ domain, GATA domain, RS domain, GAR domain</i>
Boxes, Response Elements and Sites	Name can refer to 1) the sequence or site itself (<i>TAAG</i>), 2) a non-protein that binds to it (<i>Glucocorticoid Response Element</i>), 3) a protein that binds to it (<i>Sp1</i>), or 4) to homologous genes (<i>VL30</i>).	Tag only if the sequence or site refers to a gene or protein.	<i>VL30 element, Zta response elements, activating protein 1 (AP-1) site, Ets binding site, SPI site, AP-2 box</i>	<i>GRE, TRE, cyclic AMP response element (CRE), TAAG sites, TGn motif, TAR element, UP element</i>
Hormones	Some are peptide hormones.	Tag only peptide hormones.	<i>Insulin, Glucagon, growth hormone</i>	<i>Estrogen, Progesterone, thyroid hormone</i>
“and” constructs	Some conjuncts require the entire construct.	Unless both conjuncts can stand alone, tag them together.	<i>TCR alpha and beta, D-lactate and D-glycerate dehydrogenase</i>	<i>TCR alpha, beta, D-lactate, D-glycerate dehydrogenase</i>
Viral Sequences	Promoters, enhancers, repeats are distinguished by organism.	Tag only if the organism is present.	<i>Viral LTR, HIV long terminal repeat, SV40 promoter</i>	<i>LTR, long terminal repeat</i>
Sequences	Some sequences lack gene or protein names.	Tag only if a gene name is included.	<i>NF kappa B enhancer (TGGAAATTCC)</i>	<i>TCTTAT, TTGGGG repeats</i>
Embedded Names	Some names are embedded in non-gene text.	Tag only the gene part.	<i>P-47-deficient, ras-transformed</i>	<i>P-47-deficient, ras-transformed</i>
Transposons, Satellites	Often repetitive sequences.	Tag if specific.	<i>LI element, TN44, copia retrotransposon</i>	<i>non-LTR retrotransposon</i>
Antibodies	Often use organism or disease name.	Tag if specific.	<i>anti-SF group rickettsiae (SFGR)</i>	<i>antinuclear antibody</i>
Alternative Transcripts	Names differ from primary transcript.	Tag if primary transcript named.	<i>I kappa B gamma, VEGF20</i>	<i>Exon 2, IIA</i>

Table 2: Some problematic gene/protein annotations and conventions followed in GENETAG-05.