



Comparative experiments on learning information extractors for proteins and their interactions

Razvan Bunescu^{a,1,5}, Ruifang Ge^{a,1,5}, Rohit J. Kate^{a,5},
Edward M. Marcotte^{b,1,2,5}, Raymond J. Mooney^{a,*,1,3,5},
Arun K. Ramani^{b,5}, Yuk Wah Wong^{a,4,5}

^aDepartment of Computer Sciences, University of Texas, Austin, TX 78712, USA

^bInstitute for Cellular and Molecular Biology and Center for Computational Biology and Bioinformatics, University of Texas, Austin, TX 78712, USA

Received 16 December 2002; received in revised form 14 July 2004; accepted 16 July 2004

KEYWORDS

Information extraction;
Text mining;
Machine learning;
Protein interactions;
Medline

Summary

Objective: Automatically extracting information from biomedical text holds the promise of easily consolidating large amounts of biological knowledge in computer-accessible form. This strategy is particularly attractive for extracting data relevant to genes of the human genome from the 11 million abstracts in Medline. However, extraction efforts have been frustrated by the lack of conventions for describing human genes and proteins. We have developed and evaluated a variety of learned information extraction systems for identifying human protein names in Medline abstracts and subsequently extracting information on interactions between the proteins.

Methods and Material: We used a variety of machine learning methods to automatically develop information extraction systems for extracting information on gene/protein name, function and interactions from Medline abstracts. We present cross-validated results on identifying human proteins and their interactions by training and testing on a set of approximately 1000 manually-annotated Medline abstracts that discuss human genes/proteins.

* Corresponding author. Tel: +1 512 471 9558; fax: +1 512 471 8885.

E-mail addresses: razvan@cs.utexas.edu (R. Bunescu), grf@cs.utexas.edu (R. Ge), rjkate@cs.utexas.edu (R.J. Kate), marcotte@icmb.utexas.edu (E.M. Marcotte), mooney@cs.utexas.edu (R.J. Mooney), arun@icmb.utexas.edu (A.K. Ramani), ywwong@cs.utexas.edu (Y.W. Wong)

URL: <http://www.cm.utexas.edu/faculty/Marcotte.html>, <http://www.cs.utexas.edu/users/mooney>.

¹ Supported by grant IIS-0325116 from the National Science Foundation.

² Supported by the Welch Foundation (F-1515), the National Science Foundation (ITR-0219061), and the Texas Advanced Research Program.

³ Supported by grant IIS-0117308 from the National Science Foundation.

⁴ Supported by an MCD Fellowship from the University of Texas at Austin.

⁵ Author names are in alphabetical order.

Results: We demonstrate that machine learning approaches using support vector machines and maximum entropy are able to identify human proteins with higher accuracy than several previous approaches. We also demonstrate that various rule induction methods are able to identify protein interactions with higher precision than manually-developed rules.

Conclusion: Our results show that it is promising to use machine learning to automatically build systems for extracting information from biomedical text. The results also give a broad picture of the relative strengths of a wide variety of methods when tested on a reasonably large human-annotated corpus.

© 2004 Elsevier B.V. All rights reserved.

1. Introduction

An incredible wealth of biological information generated using biochemical and genetic approaches is stored in published articles in scientific journals. Summaries of more than 11 million such articles are available in the Medline database. However, retrieving and processing this information is very difficult due to the lack of formal structure in the natural-language narrative in these documents. Automatically extracting information from biomedical text holds the promise of easily consolidating large amounts of biological knowledge in computer-accessible form. Information extraction (IE) systems could potentially gather information on global gene relationships, gene functions, protein interactions, gene-disease relationships, and other important information on biological processes.

A number of recent projects [1–10] have focused on the manual development of IE systems for extracting information from biomedical literature. Unfortunately, manual engineering of information extraction (IE) systems for particular applications is a tedious and time-consuming process [11]. Each new type of information to be extracted requires a significant new engineering effort to develop specific extraction patterns for identifying this information. Human-developed rules are also rarely able to accurately capture all of the variety of formats and contexts in which the desired information can appear in natural-language documents.

Consequently, significant recent research in information extraction has focused on using machine learning techniques to help automate the development of IE systems [12,13]. A number of machine learning methods, including grammar induction, hidden Markov models, inductive logic programming, naive Bayes text categorization, and decision tree induction, have been used to help automate the development of IE systems. First, learning systems are trained on a corpus of documents in which human experts have tagged the desired information. Next, the IE systems induced from this supervised data are used to extract new

information from novel test documents. Some projects on extracting information from biomedical literature have also employed such learning techniques [14–21].

We are exploring the use of a variety of machine learning methods to automatically develop IE systems for extracting information on gene/protein name, function and interactions from Medline abstracts. For our purposes, genes and proteins are interchangeable since, typically, there is a direct correspondence between proteins and the genes that code for them. We focus specifically on extracting information about human genes and proteins. Approximately 40,000 human genes are known from the sequences of the human genome [22,23], yet fewer than 5000 are well characterized and likely to be described in the literature. Unlike other organisms, such as yeast or *E. coli*, human gene names have no standardized naming convention, and thus represent one of the most difficult set of gene/protein names to extract. For example, human genes/proteins may be named with standard English words, such as “light”, “map”, “complement”, and “Sonic hedgehog”. Names may be alphanumeric, may include Greek or Roman letters, may be case sensitive, and may be composed of multiple words. Names are frequently substrings of each other, such as “epidermal growth factor” and “epidermal growth factor receptor”, which refer to two distinct proteins. It is therefore necessary that an information extraction algorithm be specifically trained to extract gene and protein names accurately.

In this paper, we present results on learning to extract human protein names and their interactions. We employ a variety of learning methods including pattern-matching rule induction (RAPIER) [24], boosted wrapper induction (BWI) [25], memory-based learning (MBL) [26], transformation-based learning (TBL) [27], support vector machines (SVMs) [28], and maximum entropy (MaxEnt) [29]. We present cross-validated results on identifying human proteins and their interactions by training and testing on a set of approximately 1000 manu-

```

PMID -- 9367879
TI -- A c - Cbl yeast two hybrid screen reveals interactions
with 14 - 3 - 3 isoforms and cytoskeletal components .
PG -- 46 - 50 AB - The protein product of c - cbl1,2,3 proto -
oncogene is known to interact with several proteins , including
Grb21, Crk2 and PI3 kinase3 , and is thought to regulate
signalling by many cell surface receptors .
The precise function of c - Cbl in these pathways is not clear ,
although a genetic analysis in Caenorhabditis elegans suggests
that c - Cbl4 is a negative regulator of the epidermal growth
factor receptor4 . Here we describe a yeast two hybrid screen
performed with c - Cbl in an attempt to further elucidate
its role in signal transduction . The screen identified
interactions involving c - Cbl5,6 and two 14 - 3 - 3 isoforms
, cytokeratin 185 , human unconventional myosin IC , and a
recently identified SH3 domain containing protein , SH3 P176 .
We have used the yeast two hybrid assay to localise regions of
c - Cbl required for its interaction with each of the proteins
. Interaction with 14 - 3 - 3 is demonstrated in mammalian cell
extracts .
AD -- Trescowthick Research Laboratories , Peter MacCallum
Cancer Institute .

```

Figure 1 Abstract with all the proteins and interactions tagged. The protein names have been underlined and their same subscript numbers indicate interaction between the proteins.

ally-annotated Medline abstracts that discuss human genes/proteins. Previous projects on extraction from Medline typically present results for a single method on somewhat smaller corpora with limited or no comparison to other methods. By contrast, we present uniform results of a wide variety of methods on a single, reasonably large, human-annotated corpus, thereby giving a broader picture of the relative strengths of different approaches.

2. Biomedical corpora

2.1. Tagging of medline abstracts

In order to generate a corpus of training and test data for extracting protein names and protein interactions, we manually tagged approximately 1000 abstracts (including the titles) from among the 11 million abstracts available in Medline. Tagging was performed using an existing IE-tagging tool¹ modified to enhance file handling and to retain negative examples. This program accepts a directory of files to be tagged and allows the user to tag them using a graphical interface based on a file of possible labels and writes the SGML tagged files into an output

directory. Three annotated data sets were generated:

- (1) Seven hundred fifty abstracts containing the word "human" were extracted from the Medline database and tagged for gene/protein names. 61.3% of the abstracts discussed gene/protein names, for a total of 5206 names. An example of a tagged abstract is shown in Fig. 1.
- (2) Two hundred abstracts previously known to contain protein interactions were obtained from the Database of Interacting Proteins (DIP [30]) and tagged for both 1101 protein interactions and 4141 protein names. An example is shown in Fig. 1.
- (3) As negative examples for protein interactions were rare in (2), a set of 30 abstracts were manually selected such that they had sentences with more than one gene but the abstracts did not talk about any gene interactions.

We used data set (1) for testing protein names, and data sets (2) and (3) for testing protein interactions.

2.2. Rules used for tagging

Due to the ambiguities involved in human gene/protein names and interactions, it was necessary

¹ URL:http://www-2.cs.cmu.edu/~kseymore/general_tagger.pl.

to develop a set of conventions for their consistent tagging. In the following discussion, we indicate protein names by underlined text and their same subscript numbers indicate interaction between the proteins. Manual examination of many abstracts revealed several ambiguities, such as whether the organism names should be tagged (e.g. human delta catenin or human delta catenin), whether punctuation should be tagged (e.g. (LIGHT) or (LIGHT)), and whether generic protein family names should be tagged (e.g. armadillo protein p0071 or armadillo protein p0071). Such cases led to the following set of tagging conventions:

- (1) As few extra characters as possible are tagged. Punctuation marks and plural characters are not tagged.
- (2) Gene/protein names are tagged regardless of context, even when gene names are substrings of other gene names. (e.g. GITR ligand)
- (3) Generic protein/gene families are not tagged, only specific names which could ultimately be traced back to specific genes in the human genome. (e.g. “Tumor necrosis factor” would not be tagged, while “tumor necrosis factor alpha” would be.)
- (4) Tags for interacting proteins follow the same conventions as for other proteins. All stated instances of protein interactions are tagged, even when tags are nested. (e.g. human GITR₁ligand (hGITR₁))

3. Protein name identification

Named entity recognition (NER), identifying names of people, organizations, and places in text, is a well studied problem in information extraction from news articles. In recent years, machine learning approaches have become the standard in developing robust, accurate NER systems [31,32]. Biomedical applications have special types of named entities that are different from those typically addressed by existing NER systems. These include names of diseases, genes, proteins, organisms, organs, organelles, and other biological entities. In this section, we explore the problem of recognizing references to human proteins using the tagged data described in the previous section.

3.1. IE methods

3.1.1. Dictionary-based extraction

The success of a protein tagger depends on how well it captures the regularities of protein naming as well as name variations. In the dictionary-based approach, we started with an extensive set of pro-

<i>Protein name (OD)</i>	<i>Generalized name (GD)</i>	<i>Canonical form (CD)</i>
interleukin-1 beta	interleukin $\langle n \rangle \langle g \rangle$	interleukin
interferon alpha-D	interferon $\langle g \rangle \langle r \rangle$	interferon
NF-IL6-beta	NF IL $\langle n \rangle \langle g \rangle$	NF IL
TR2	TR $\langle n \rangle$	TR
NF-kappa B	NF $\langle g \rangle \langle r \rangle$	NF

Figure 2 Dictionary generalizations.

tein names extracted from two fairly comprehensive sources:

- (1) The file `human.seq`, downloaded from the Human Proteome Initiative (HPI) of EXPASY.²
- (2) The file `feb2002-tables.tar.gz`, downloaded from the Gene Ontology Database.³

Altogether, these dictionaries contain 42,172 gene/protein names (synonyms included). This collection of protein names, henceforth referred to as the original dictionary (OD), was further extended using a generalization procedure to obtain a generalized dictionary (GD). The aim was to extend the coverage of the original set, while at the same time trying to minimize any decrease in accuracy.

Generalizing a dictionary entry involved identifying those parts susceptible to change in new protein names, and replacing them with generic placeholders. Thus, we isolate and replace numbers with $\langle n \rangle$, Roman letters with $\langle r \rangle$ and Greek letters with $\langle g \rangle$. Fig. 2 shows some examples of name generalizations.

In the GD-based extraction, we tag a textual n -gram as a protein name only if it is an instance of one of the generalizations from the generic dictionary. To extend the coverage even more, we have created a canonical dictionary (CD) consisting of canonical forms of protein names. A canonical form is obtained from a generic form by stripping it of all generic tags, as can be seen in the examples from Fig. 2. From the resulting set we filter out common English words whose presence could lead to a decrease in accuracy. Consequently, in the CD-based extraction, a textual n -gram is deemed as being a protein name if its canonical form is part of the canonical dictionary. Both GD and CD introduce spurious entries in the dictionary, leading to a decrease in precision. For instance, because “HT3” is an entry in OD, its generalization $HT\langle n \rangle$ will cover “HT 29”, which is a cell line. Also, CD will match words that are classes of proteins, and not particular proteins, as is the case with “oncogene” which was derived as the canonical form of “oncogene 24P3”. Because the tagging based on both the

² URL: http://us.expasy.org/sprot/hpi/hpi_ftp.html.

³ URL: <http://www.godatabase.org/dev/database/archive>.

Table 1 Performance of protein taggers in various settings

IE methods and additional information used	Precision(%)	Recall(%)	F-measure(%)
Dictionary-based			
Original dictionary	56.70	27.24	36.80
Plus generalized dictionary	62.27	45.85	52.81
Plus canonical dictionary	41.88	54.42	47.33
RAPIER			
Words only	76.11	9.97	17.63
Part-of-speech	70.84	11.05	19.12
Dictionary-based tagger	74.49	12.22	21.00
BWI (300 iterations, 2 lookaheads, max. recall)			
Words only	70.67	11.52	19.81
Dictionary-based tagger	71.01	24.06	35.94
k -NN ($k = 1, N = 2$)			
Part-of-speech	34.66	40.66	37.42
Dictionary-based tagger	47.30	47.82	47.56
TBL			
Words only	47.08	36.65	41.22
Dictionary-based tagger	56.80	34.62	43.02
SVM ($N = 2$, full training set, max. recall)			
Preceding class labels	69.16	19.74	30.72
Preceding class labels and part-of-speech	70.18	19.72	30.79
Preceding class labels and dictionary-based tagger	65.00	45.43	53.48
with additional suffix features	70.38	44.49	54.42
MaxEnt ($N = 1$, Viterbi w/o greedy extraction, max. recall)			
W/o dictionary	71.10	42.31	53.05
With dictionary	73.37	47.76	57.86
With dictionary, two tags only (I,O)	66.41	44.74	53.46
KEX	14.68	31.83	20.09
ABGENE	32.39	45.87	37.97

original and generic dictionary gave better results than other combinations (as shown in the first entry of Table 1), we used this particular dictionary-based tagger for supplying a pre-tagged input to some of the learning methods that will be discussed in the following sections.

3.1.2. RAPIER

RAPIER[24] is a rule learning algorithm that acquires unbounded patterns for extracting information from text. Each extraction rule consists of three parts: (1) a pre-filler pattern that matches text immediately preceding a filler (e.g. a protein name), (2) a filler pattern that matches the extracted substring, and (3) a post-filler pattern that matches the text immediately following the filler. RAPIER begins with a most-specific set of rules and compresses the rule base by repeatedly replacing rules with more general ones.

To construct the initial rule base, most-specific patterns are created for each training example, specifying words for the filler, all words in the text preceding the filler, and all words in the text following the filler. To generate new rules, pairs of existing

rules are randomly selected and their least-general generalizations created. RAPIER starts with rules containing only generalizations of the filler patterns, and uses beam search to efficiently specialize the rules by adding pieces of the generalizations of the pre- and post-filler patterns of the seed rules, until the best rule in terms of information gain produces no spurious fillers when matched against the training examples. The best generalized rule is then added to the rule base, and the process repeats until compression has failed more than a specified number of times.

To help RAPIER capture generalities that are not evident from the words alone, we supplied additional syntactic and semantic information to the learner in some of our experiments. First, we added part-of-speech (POS) tags to every word in the text. POS tags are potentially useful because certain types of words (e.g. cardinal numbers and proper nouns) are likely candidates of being parts of a protein name.

In another experiment, we included the output of the dictionary-based tagger (Section 3.1.1) in place

```

PMID -- 11529898
Lol p 1 is one of the most important allergens in grass pollen
extracts...
Lol p 1 is one of the most important allergens in grass pollen
extracts...
BBBB Lol EEEE BBBB p 1 EEEE is one of the most important
allergens in grass pollen extracts...

```

Figure 3 Incorporating information from the dictionary-based tagger. The first sentence contains the correct tagging. The second sentence is the output of the dictionary-based tagger. The third sentence shows the input for RAPIER and BWI. The output tags of dictionary-based tagger have been transformed into special tokens BBBB and EEEE standing for begin and end of the tags respectively.

of the POS tags in the form of special tokens (see Fig. 3). By adding these tokens, we incorporated domain knowledge into the learning algorithm. At the same time, the learning algorithm can find general patterns that refine the output of the dictionary-based tagger.

3.1.3. Boosted wrapper induction

Boosted wrapper induction (BWI) [25] learns extraction rules composed only of simple contextual patterns called *wrappers* [33]. Although wrappers are highly accurate predictors of the start or end of a protein name, each of them has limited coverage since Medline abstracts do not exhibit a rigid structure. BWI circumvents this limitation by using boosting [34], which repeatedly learns simple, weak patterns that focus on the training examples for which the previous patterns have done poorly. The predictions of all learned patterns are then combined using a weighted voting scheme. The result is a boosted wrapper, which has been shown to be successful in several natural text domains.

To perform protein-name extraction using a boosted wrapper, every word boundary i in a Medline abstract is first given a *fore* score $F(i)$, which indicates its likelihood of being the start of a protein name, and an *aft* score $A(i)$, which indicates its likelihood of being the end of a protein name. Then, the wrapper recognizes a text fragment (i, j) as a protein name if and only if $F(i)A(j)H(j-i) > \tau$, where $H(k)$ is a function that reflects the probability that a protein name has length k , and τ is a numeric threshold that controls the level of recall. By varying τ , we are able to perform extraction at different degrees of confidence.

In our experiments with BWI, we tested the usefulness of including the output of the dictionary-based tagger (Section 3.1.1) as part of the input of the learner, in the same way as it was done in Section 3.1.2.

3.1.4. Support vector machines

Support Vector Machines (SVMs) are one of the most recently developed classification methods [35]. They are well-founded in computational learning theory, and have been shown to generalize well in the presence of very many features. They are generally considered to be the currently best technique for text classification [36].

Assume that all m training examples consist of a vector of n features, and belong to either positive or negative class as follows: $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_m, y_m)$, where $\mathbf{x}_i \in \mathbf{R}^n$ is the i -th feature vector and $y_i \in \{+1, -1\}$ is its class label. Then an SVM learns an optimal threshold function $f(\mathbf{x}) = \langle \mathbf{w}, \mathbf{x} \rangle + b$, $\mathbf{w} \in \mathbf{R}^n, b \in \mathbf{R}$, which separates the training examples into two classes. An example \mathbf{x} is classified as positive when $f(\mathbf{x}) > 0$, or negative when $f(\mathbf{x}) < 0$. A threshold function is optimal when the margin of separation between the two classes is maximal. It can be proven that the margin is maximized when the norm of \mathbf{w} is minimized. This leads to a constrained quadratic optimization problem which can be exactly solved efficiently.

Since our tagged Medline abstracts do not contain any protein names that directly abut each other, we can reduce the NER problem to classification of individual words. First, an SVM classifier determines if each word is part of a protein name or not, by looking at the word itself and its surrounding context. Next, protein names are extracted by identifying the longest sequences of words that have been classified as parts of a protein name. Similar approaches have been applied successfully to the task of *text chunking*, which is identifying simple phrases such as non-recursive noun and verb phrases [37,32].

For each token, we built a feature vector consisting of the current word, the previous and the following N words. We also included POS tags generated by the Brill's tagger⁴ and the output of the

⁴ URL: http://www.cs.jhu.edu/~brill/RBT1_14.tar.Z.

dictionary-based protein tagger (Section 3.1.1) for all $2N + 1$ words. We ignored capitalization when preparing the feature vectors to avoid sparsity. To capture morphological similarities and alleviate the problem of unseen words, we included as features the last one, two, and three characters of each word in the feature vector, which we henceforth refer to as the *suffix features*. Inspired by the text chunking algorithm presented in [38], we included the class labels of the two preceding words as part of the feature vector. Since the class labels were not given in the test data, they were decided dynamically during the tagging of previous words. Because numerical values were needed, each word or tag in each position was a separate binary feature. For each extracted sequence of tokens, we used the minimal distance from the hyperplane $f(\mathbf{x}) = 0$ as a quantitative measure of confidence. For the inner product $\langle \mathbf{w}, \mathbf{x} \rangle$, we used $\mathbf{w}^T \mathbf{x}$, which resulted in a linear threshold function. It has been argued that most text categorization problems are linearly separable [36], so in our case a linear threshold function should suffice. We used version 5.0 of SVM^{light},⁵ which is highly efficient in dealing with sparse instances.

The training set for the token classification problem is highly imbalanced. Out of the 209,022 tokens in our corpus, only 10,175 of them (4.87%) are protein names. As pointed out by [39], the induced classifiers tend to be highly accurate on negative examples but also produce many false negatives which lead to low recall. By sampling the training set and feeding the learner with only negative examples surrounding the positive ones, we can shift the resulting hyperplane and potentially reduce the number of false negatives. Our experiments supported this claim and showed that we could attain very high recall at the expense of precision.

3.1.5. Maximum entropy

Maximum Entropy [29] is a widely used method for inducing probabilistic classifiers. The classification problem is viewed in terms of a random process that produces an output value y from a finite set Y , based on a contextual information x , a member of a finite set X . In a tagging scenario, this means associating a tag y to each text token, whereas the context x can be derived from the text centered at the current token position. In maximum entropy modeling we are looking for a probability distribution $p(y|x)$

expressed in terms of a set of user specified features $f_i(x, y) \in F$:

$$p(y|x) = \frac{1}{Z(x)} \exp\left(\sum_i \lambda_i f_i(x, y)\right)$$

where $Z(x) = \sum_y \exp(\sum_i \lambda_i f_i(x, y))$ is a normalizing constant. Each feature f_i is a binary function based on the current context x and its proposed classification y .

In the case of maximum entropy tagging (henceforth referred to as MaxEnt), we distinguish among five types of tags in Y (as opposed to using only two tags, as was the case with SVMs):

- *S(-tart)*- indicates the first token of a protein name.
- *E(-nd)*- indicates the last token of a protein name.
- *C(-ontinue)*- indicates a token strictly inside a protein name.
- *U(-nique)*- indicates the unique token of a protein name.
- *O(-ther)*- all other tokens (outside protein names).

We hypothesize that the task of tagging the first, the last, or the unique token of a protein name is slightly different from that of tagging other tokens inside a protein name, hence the extended set of tags.

The abstracts are tokenized, segmented in sentences, and annotated with part-of-speech tags using the same tools as in Section 3.1.4. Then the model generates feature vectors by scanning each pair (x_i, y_i) in the training data using the feature templates given in Fig. 4. We use a threshold of three as the minimum number of times that a feature should appear in the training data in order to be considered. The word class features *cf* and *cb* are based on the similar features introduced in [40]. Thus, for a character x we define *type*(x) as ‘A’ if x is an upper-case letter, ‘a’ if x is a lower-case letter, ‘0’ if x is a digit and x otherwise. The *cf* feature then is

Name	Feature description	Feature body
<i>w</i>	current word	$w(x_i) = \langle w \rangle$ & $y_i = \langle y \rangle$
<i>pw</i>	previous word	$pw(x_i) = \langle w \rangle$ & $y_i = \langle y \rangle$
<i>nw</i>	next word	$nw(x_i) = \langle w \rangle$ & $y_i = \langle y \rangle$
<i>pos</i>	POS, current word	$pos(x_i) = \langle pos \rangle$ & $y_i = \langle y \rangle$
<i>ppos</i>	POS, previous word	$ppos(x_i) = \langle pos \rangle$ & $y_i = \langle y \rangle$
<i>npos</i>	POS, next word	$npos(x_i) = \langle pos \rangle$ & $y_i = \langle y \rangle$
<i>cf</i>	word class (full)	$cf(x_i) = \langle cf \rangle$ & $y_i = \langle y \rangle$
<i>cb</i>	word class (brief)	$cb(x_i) = \langle cb \rangle$ & $y_i = \langle y \rangle$
<i>dict</i>	dictionary tag	$dict(x_i) = \langle dt \rangle$ & $y_i = \langle y \rangle$
<i>pt</i>	previous tag	$pt(x_i) = \langle y' \rangle$ & $y_i = \langle y \rangle$

Figure 4 Feature templates.

⁵ URL: <http://svmlight.joachims.org/>.

$conf(W, u, v)$	
Input:	W , a sequence of tokens w_1, w_2, \dots, w_T $[u, v]$, an extraction span with $1 \leq u \leq v \leq T$
Use the forward procedure on W with $p(y x)$ to compute: $\alpha_t(y) = p(y_t = y W)$, where $1 \leq t \leq T$ and $y \in \{S, C, E, U, O\}$	
if $u = v$ $conf_{avg}(W, u, v) = conf_{min}(W, u, v) = \alpha_u(U)$	
else $p_u = \alpha_u(S)$ $p_{u+1} = p(y_{u+1} = C W, y_u = S)$... $p_v = p(y_v = E W, y_{v-1} = C)$ $conf_{avg}(W, u, v) = \frac{1}{v-u+1} \sum_{t=u}^v p_t$ $conf_{min}(W, u, v) = \min_{u \leq t \leq v} (p_t)$	

Figure 5 Extraction confidence.

the current word with each character mapped to its type, while the brief version *bf* results from *cf* by removing repeating character types. For example, if “FGF1” is the current word, then *cf*=‘AAA0’, and *bf* = ‘A0’. Another special feature is *pt*, based on the tag assigned to the previous token. The dependence of the current tagging decision on the previous tag, unknown during testing, forces us to consider all possible tags for the previous token when tagging unseen data. For a particular token sequence (tokenized sentence), this will result in a potentially very large set of possible taggings. The classical approach is to use a Viterbi-like algorithm for finding the most likely sequence of tags. To each of the resulting extractions, we associate a confidence measure. The exact calculation of the two confidence measures $conf_{avg}$ and $conf_{min}$ is described in Fig. 5. The Viterbi and forward procedures used therein are similar with those used for hidden Markov models [41]. Varying the confidence level will later allow us to trade off between precision and recall (see Section 3.2.2).

We chose to take either the minimum or the average as we were targeting a length-independent measure. Also, when using the average function, or the minimum function, one has to ensure that the quantities involved have a similar interpretation and consequently can be safely combined. In our case, we can view the value $\alpha_t(y)$ as another transition probability, namely the probability of reaching state y at time step t from a special state encoding the beginning of the sentence.

One drawback of using the Viterbi algorithm is that by focusing on the most likely sequence of tags, the program is missing many low confidence extractions that might help in extending the recall endpoint. When applied on test data, the Viterbi algorithm, augmented with the confidence measure $conf_{min}$, results in a maximum recall of 47.76%. To further extend it, we use the greedy algorithm from

$greedy_extract(W, u, v)$	
Input:	W , a sequence of tokens w_1, w_2, \dots, w_T $[u, v]$, an extraction domain with $1 \leq u \leq v \leq T$
if $u > v$ return \emptyset	
else $[l, r] = \operatorname{argmax}_{[l,r] \subseteq [u,v]} conf(W, l, r)$ $LE = greedy_extract(W, u, l-1)$ $RE = greedy_extract(W, r+1, v)$ return $LE \cup \{[l, r]\} \cup RE$	

Figure 6 Greedy extraction.

Fig. 6 on all token sequences appearing between two consecutive Viterbi extractions, thus obtaining additional extractions compatible with the set of proteins already extracted through the Viterbi procedure (two extractions are compatible if they do not overlap). All the results that presuppose using a confidence measure are based on $conf_{min}$, which does a better job at extending the recall endpoint.

We base our maximum entropy approach on the `opennlp.maxent` package,⁶ version 2.1.0, which uses the generalized iterative scaling algorithm [42] for estimating the parameters of the log-linear model.

3.1.6. Existing protein name identification systems

We also tested two existing protein name identification systems. The first one is KEX version 1.21, which is based on the PROPER algorithm described in [1]. It consists of a set of hand-built pattern matching rules which makes use of part-of-speech information given by the Brill’s tagger. Without depending on any protein-name dictionaries, KEX has been reported to achieve 94.70% precision and 98.84% recall on a corpus of 80 abstracts on SH3 and signal transduction domains.

The second system is ABGENE, introduced in [43]. ABGENE uses a transformation-based tagger to produce an initial tagging. Then it employs a number of dictionaries and contextual rules to weed out false positive and recover false negative. It was tested on a corpus consisting of the complete set of abstracts introduced into Medline between 15 June and 24 September 2001, and was reported to give good results.

3.2. Experimental results

We begin this section by explaining the methodology followed in our experiments. We present next the quantitative results of the IE methods used for

⁶ URL: <http://maxent.sourceforge.net/>.

extracting protein names. The section ends with a comparative analysis of the results.

3.2.1. Experimental methodology

The 750 Medline abstracts annotated with protein tags were tokenized using simple pattern rules developed for the Penn Treebank project [44]. For programs requiring sentence-segmented input, we used the sentence segmenter from the KEX tagger with additional rules for bulleted lists. For those learning algorithms requiring POS tags, we used Brill's POS tagger, which we trained by using 10000 untagged Medline abstracts as the training set. Those abstracts were obtained the same way we did for the 750 abstracts. No stemming or stopword filtering was performed during the experiments. Capitalization was retained unless otherwise specified.

We performed 10-fold cross validation on each learning algorithm with a particular parameter setting. This provides average performance over ten random trials, each training on 90% of the data and testing on the remaining 10%. Each extracted protein name in the test data was compared to the human-tagged data, with the positions taken into account. Since ABGENE provides no positional information, we assume that all occurrences of its extracted strings are recognized as protein names. Two protein names are considered a match if they consist of the same character sequence in the same position in the text. This detects circumstances where common English words are incorrectly recognized as protein names (e.g. "light", "at"), and ensures that all references to each protein are recognized. We measured precision (percentage of extracted names that are correct), recall (percentage of correct names that are found), and *F*-measure (harmonic mean of precision and recall) [11].

3.2.2. Quantitative results

Table 1 summarizes results for the protein taggers presented in Section 3.1, along with any additional sources of information used. We also include results obtained with two additional taggers: one using transformation-based learning (TBL) [27], and another based on the *k*-nearest neighbor (*k*-NN) method in which classification is done by extrapolation from the *k* most similar training examples. For systems that output confidences that allow trading-off precision and recall (i.e. BWI, *k*-NN, SVM and MaxEnt), results are presented for the maximum achievable recall or the best *F*-measure.

For ease of comparison, we show recall–precision curves in Fig. 7, using the version of each system that gave the best *F*-measure (as shown in bold in

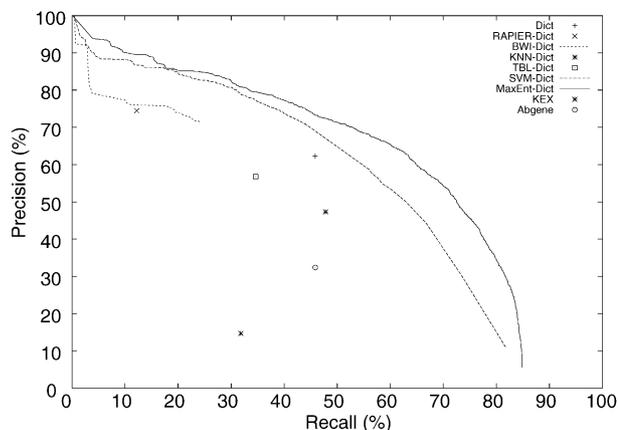


Figure 7 Precision-recall curves for protein taggers on the 750 abstracts dataset.

Table 1). For those IE methods that output extraction confidences, we show curves indicating the precision for each achievable level of recall. Single recall–precision points are shown for all other methods.

Given that the MaxEnt approach achieves the best results on the 750 abstracts dataset, we applied it on the 230 abstracts from the interactions dataset, our aim being to feed these automatically tagged abstracts to the interaction extraction program (see Section 4.2.2 for overall results of the combined approach). The tagging performance on the interaction dataset is shown in Fig. 8(a).

3.2.3. Discussion of results

Overall, the results show limited utility of POS tags. The use of POS tags in RAPIER, *k*-NN, and SVM does not improve *F*-measure significantly according to a paired *t*-test ($p > 0.05$). While the dictionary-based tagger barely improves *F*-measure for RAPIER and TBL, it is useful for the rest of the learning methods to different extents. It improves both precision and recall for BWI, *k*-NN and MaxEnt, while for SVM it hurts the precision slightly, but this is outweighed by a larger gain in recall.

Out of all the learning methods tested here, SVM and MaxEnt achieve a significant improvement over the dictionary-based tagger in terms of *F*-measure. Another advantage is that both are able to achieve arbitrarily high precision by adjusting the confidence level. This is possible because the extraction confidence is highly correlated with the probability of correctness. Since high precision is needed to extract accurate knowledge from text, this is a significant contribution.

We have also included results for MaxEnt with a tagging scheme based on two tags only. The difference in performance validates our initial hypothesis

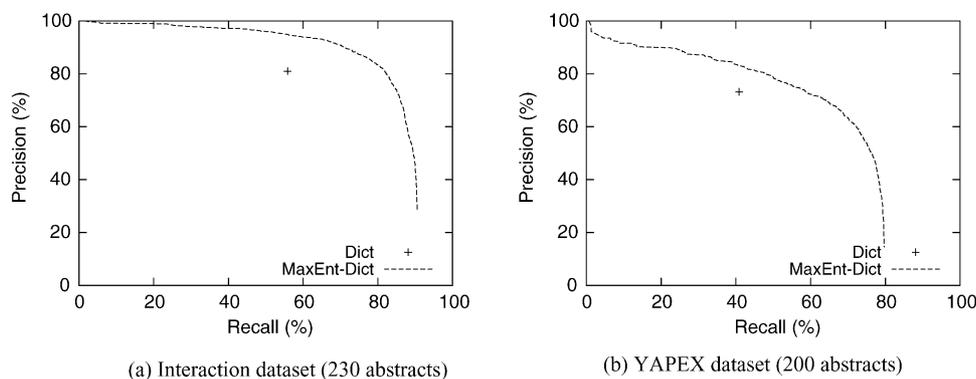


Figure 8 Performance of MaxEnt protein tagger on two other datasets.

suggesting the use of more than one tag for tokens inside a protein name.

All of our IE methods perform significantly better than two existing protein taggers, KEX and ABGENE. Given that these systems were developed for different distributions of proteins, this is not surprising; however it does illustrate the relative difficulty of identifying human proteins. The hand-built rules used in KEX were developed and tested on a rather confined set of proteins different from the human proteins in our data. ABGENE uses a version of TBL to learn a protein tagger; however, the specific tagger we obtained was not trained specifically for human proteins. Our own TBL system is more indicative of the performance of this approach when specifically trained for human proteins; however, note that many of the other learning approaches perform better than TBL.

As shown in Fig. 8(a), the performance on the interaction dataset is a lot better than on the protein dataset, and this is also reflected in the results of the dictionary-based tagger. There are two main reasons for this significant difference:

- (1) The protein dataset has been manually tagged by 9 people, in just one pass. After analyzing the results, we have discovered significant tagging inconsistencies which clearly affected the learning performance. On the other hand, the interaction dataset has been tagged by one person only, resulting in a more consistent tagging.
- (2) Each of the 230 interaction abstracts contains at least two proteins, due to the particular selection process described in Section 2.1. This results in a significant bias which is captured by the learning algorithm. Comparatively, 38.7% of the 750 abstracts dataset contain no proteins, while many of the same abstracts include various names for cell lines, or amino acids, names which are very similar with protein names, mak-

ing the task of recognizing proteins more realistic, but at the same time harder.

We have also tried our protein name extraction systems on the YapeX⁷ dataset (200 Medline abstracts) and the results (Fig. 8(b)) are comparable with those obtained on our corpus of 750 abstracts (Fig. 7).

4. Protein interaction extraction

Identifying *relations* between named entities stated in text is a more difficult IE problem that only recently has attracted significant attention in research on extraction from new articles. The current ACE (Automated Content Extraction) program at the National Institute of Standards and Technology (NIST) [45] is focused on identifying various social, action-role, part-of, and locational relations between named entities. Several projects have focused on extracting relations from biomedical text, such as identifying gene-disease relations, subcellular localizations, or protein interactions [15,17,18,2–6,46]. This section discusses our work on identifying human-protein interactions assuming that the proteins themselves have already been tagged, and shows that machine-learning systems out-perform human-written extraction rules with respect to providing a wider range of precision and recall.

4.1. IE methods

4.1.1. RAPIER and boosted wrapper induction

In order to adapt slot-filling IE systems that extract individual entities (like RAPIER and BWI) to the problem of extracting *relations*, we developed two approaches. The first approach we call the *Inter-*

⁷ URL: <http://www.sics.se/humle/projects/prothalt/>.

Figure 9 Interactor, interactee and interfiller.

filler approach. Given two tagged entities participating in a relationship, the text fragment between them is called the interfiller (see Fig. 9). If a slot-filling IE system extracts an interfiller, the tagged entities before and after it can be extracted as participating in the targeted relation.

The second approach we call the *Role-filler* approach. In this approach, we extract the two related entities independently into different role-specific slots. For protein interactions, we named the roles *interactor* and *interactee* (see Fig. 9). There might be many interactors and interactees extracted in one sentence, we then decide which of them participate in a relationship using the following heuristics, assuming that all interacting proteins appear in the same sentence. (1) The interactors and interactees appearing in the same sentence form a sequence of role fillers. This sequence is separated into segments at the points where an interactee is immediately followed by an interactor. Interactors and interactees can only be paired within the same segment. (2) Each interactor is associated with the next occurring interactee in the segment. (3) If there are fewer interactors (interactees) than interactees (interactors) in the segment, use the last interactor (interactee) in constructing the remaining pairs. In our human-tagged interaction corpus, assuming interactors and interactees are properly tagged, this approach identifies all the interacting pairs with 99.2% accuracy.

Both of these approaches have been used to train BWI (Section 3.1.3) to extract interacting proteins, and the Role-filler approach has been used to train RAPIER (Section 3.1.2) to extract interactions. RAPIER could not learn to extract interfillers successfully, since, in the worst case, the time complexity of its generalization algorithm can grow exponentially in the length of a filler. Since extracted entities are usually fairly short, this is typically not a problem in standard slot-filling IE. However, the long interfillers in many protein interactions prevented us from running RAPIER with the Interfiller approach.

4.1.2. Extraction using longest common subsequences (ELCS)

We have also developed a new method for directly learning patterns for extracting relations between previously tagged entities. Blaschke et al. [3,47] manually developed rules for extracting interacting proteins. Each of their rules (or frames) is a sequence of words (or POS tags) and two protein-name tokens. Between every two adjacent words is a number indicating the maximum number of intervening words allowed when matching the rule to a sentence. Here we describe a new method ELCS (Extraction using Longest Common Subsequences) that automatically learns such rules.

ELCS' rule representation is similar to that in [3,47], except that it currently does not use POS tags, but allows disjunctions of words. Fig. 10 shows an example of a rule learned by ELCS. Words in square brackets separated by ' | ' indicate disjunctive lexical constraints, i.e. one of the given words must match the sentence at that position. The numbers in parentheses between adjacent constraints indicate the maximum number of unconstrained words allowed between the two (called a *word gap*). A sentence matches the rule if and only if it satisfies the word constraints in the given order and respects the respective word gaps.

A sentence in the training data may contain more than two proteins and more than one pair of interacting proteins. In order to extract the interacting pairs, the rules should be trained to pick out exactly the interacting proteins from the sentences. To do this we replicate the sentences having n proteins ($n > 2$) into C_2^n sentences such that each one has exactly two of the proteins tagged, with the rest of the protein tags omitted. If the tagged proteins interact, then the replicated sentence is added to the set of positive sentences, otherwise it is added to the set of negative sentences. During testing, a sentence having n proteins ($n > 2$) is again replicated into C_2^n sentences in a similar way. If such a replicated sentence matches one of the rules, then the system extracts the two

Figure 10 Sample protein-extraction rule learned by ELCS. Token PROT stands for protein name.

proteins tagged in that sentence as interacting proteins.

ELCS induces rules using a bottom-up approach. Rule induction starts with maximally specific rules for each positive sentence which contain all the words in the sentence with zero-length word gaps. These are then repeatedly generalized to form more general rules until the rules become overly general and start matching negative sentences. We have developed three methods for generalizing rules. The first simple method to produce a generalization of two rules is to find the *longest common subsequence* (LCS) of words between them. Efficient algorithms for computing an LCS are presented in [48,49]. After finding the LCS between two rules, we determine the size of word gaps between every two adjacent words in their LCS as the larger of the number of words plus the sum of existing word gaps between the two LCS words where they are found in the original two rules.

Our second approach to generalization uses *edit distance* (ED) [49] and creates more specific rules that contain disjunctive constraints. The most common edit distance is Levenshtein distance [50], defined as the minimum number of edit operations (adding, deleting, or replacing an item) required to convert one sequence into another. We use the minimal edit-operation sequence obtained when computing Levenshtein distance to generalize two rules. We preserve the common word constraints between the rules, make disjunctions of constraints when one item is replaced by another in the edit sequence, and drop constraints that are added or deleted in the edit sequence. Finally, we introduce

word gaps using the method described for the LCS-based generalization.

The third generalization method finds all common sequences between the two rules and considers their *conjunction* (CJ) as the generalization. Unlike the previous two methods, this method is associative, i.e. we get the same generalization of a set of rules irrespective of the order in which we generalize two of them at a time. If there is any common pattern among the base rules then this property guarantees that the pattern will also appear in the generalization (note that it is possible to lose such a common pattern while taking LCS of two rules at a time). Word gaps are then introduced as in the previous two methods. Fig. 11 shows generalization of two sentences obtained by each of these methods.

Using one of these generalization methods, a greedy-covering, bottom-up rule-induction method is used to learn a small set of rules that cover all the positive sentences without covering many negative ones. We use an algorithm similar to beam search and consider only the r best rules for generalization at any time. We start with r randomly selected positive examples. These r rules are generalized with one of the remaining positive examples to obtain r more rules. Out of these $2r$ rules we select r rules with the highest confidence level and allow further generalization with the remaining positive examples. After iterating over the remaining positive examples in this way, the r best rules are finally included in the set of learned rules and the positive examples covered by them are removed. The entire process is repeated till we exhaust the set of positive examples.

Sentence 1: The self - association site appears to be formed by interactions between helices 1 and 2 of beta spectrin₁, repeat 17 of one dimer with helix 3 of alpha spectrin₁, repeat 1 of the other dimer to form two combined alpha - beta triple - helical segments .

Sentence 2: Title - Physical and functional interactions between the transcriptional inhibitors Id3₂ and ITF - 2b₂ .

Generalization using longest common sequence (LCS):
- (7) interactions (0) between (5) PROT (9) PROT (17) .

Generalization using edit-distance (ED):
[self|Title] (0) - (4) [be|Physical] (0) [formed|and] (0) [by|functional] (0) interactions (0) between (2) [and|the] (0) [2|transcriptional] (0) [of|inhibitors] (0) PROT (8) [of|and] (0) PROT (17) .

Generalization using conjunctions (CJ):
{ - (7) interactions (0) between (5) PROT (9) PROT (17) . } \wedge { - (11) and (6) PROT (9) PROT (17) . }

Figure 11 Generalizations of two sentences using different methods. Protein names have been underlined and same subscript numbers indicate interactions between them. Token ‘PROT’ stands for protein name.

We measure the confidence levels of our rules using m -estimate [51] which is a measure of expected accuracy of a rule. It is defined as: $confidence\ level\ (rule) = (p + m.p^+) / (p + n + m)$, where p and n are the number of positive and negative examples covered by the rule, p^+ is the prior probability of positive examples and m a parameter which should be set according to the amount of noise in the data. We set p^+ as the fraction of examples in the training data which are positive and set m based on pilot studies.

The generalizations obtained using any of the methods may result in rules that do not contain two protein-name tokens. This is fine for extracting protein interactions because we always apply the rules to sentences containing exactly two protein names (if they contain more than two protein names then we replicate the sentence as described earlier). However, constraining learned rules to contain two protein names is a useful bias. Therefore, we divide each of the training sentences in three parts: the portion of the sentence before the first protein name, the portion between the two protein names, and the portion after the second protein name. When we generalize two rules, we generalize these three parts separately. This ensures the rule will always contain two protein-name tokens. Fig. 12 shows some sample rules learned by ELCS.

4.2. Experimental results

4.2.1. Experimental methodology

Medline abstracts were pre-processed as described in Section 3.2.1. All our systems for extracting interactions require sentence segmentation since only the proteins within a sentence are considered when identifying interactions. This constraint is

satisfied by all interactions in our corpus because while manually tagging the corpus with interactions we did not find a single instance where the two interacting proteins were in different sentences. We also compared our systems with Blaschke et al.'s manually-written rules [47]. Since these rules require POS tags, we used Brill's POS tagger. We also tested a version of the human-written rules in which the POS tags are replaced by typical words indicating interactions such as *activation*, phosphorylation or *interaction* for nouns and *activates*, *binds* or *phosphorylates* for verbs, similar to the approach in [3]. These manually-written rules were developed to capture the common ways of expressing protein interactions in natural language, like "protein A binds/interacts... (with) protein B" (see the SUISEKI system overview in [47]). Such natural language constructions are general enough and not sensitive to any particular dataset. Hence, we applied these rules on our dataset to make a direct comparison with our systems.

We did two experiments to evaluate the performance of protein interaction extraction. In both experiments the machine learning systems were trained using the manually tagged abstracts (see Section 2.1) with proteins and their interactions. The two experiments differ in the way we tested the systems. In the first experiment we provide manually tagged protein names to our systems and extract interactions among these proteins. This way we get a measure of how the protein interaction extraction systems alone perform independent of the protein name extraction systems. In the second experiment we first find protein names in the abstracts using our best system for protein name extraction, MaxEnt (see Section 3.2.2), and then extract interactions among these proteins. This

```
interactions (0) between (4) PROT (0) and (4) PROT (16) .

PROT (0) / (0) PROT (10) heterodimers (36) .

[binding | substitution | AB | addition | Interestingly | TI | interactions] (0)
[of | - | ,] (3) PROT (19) [to | for | : | same | with] (10) PROT (30)
[nM | binding | 1 | CDK6 | CCR8 | death] (9) .

[linker | TI | armadillo | b558 | of] (0) [- | , | a] (5) PROT (13) [and
| / | with | to | containing] (0) PROT (2) .

{, (11) PROT (25) and (8) to (9) PROT (66) .}  $\wedge$  {, (11) PROT (16) bind
(18) PROT (66) .}

{, (10) PROT (5) for (7) PROT (9) .}  $\wedge$  {, (10) PROT (4) binding (6) PROT
(9) .}
```

Figure 12 Some example rules learned by ELCS; the first two were learned using LCS generalization, the next two using ED generalization and the last two using CJ generalization.

gives a true measure of how our systems can perform at extracting protein interactions from completely untagged abstracts. Blaschke et al.'s manually-written rules also require protein names, in this experiment we also test those rules by providing them with our extracted protein names. For this experiment, we chose the point on the MaxEnt's precision-recall curve (Fig. 8(a)) which gives 70% precision and about 90% recall for protein name extraction.

As in Section 3.2.1, performance is evaluated using 10-fold cross validation and measuring recall and precision. We consider an extracted interaction from an abstract correct only if both its proteins have been human-tagged as interacting with each other somewhere in that abstract. As the task of interest is only to extract interacting protein-pairs, in our evaluation we do not consider matching exact positions and every occurrence of interacting protein-pairs within the abstract. For those IE methods which output extraction confidences, if we extract more than one occurrence of interaction between two proteins then we combine their extraction confidences using the standard Noisy-Or method [52].

4.2.2. Quantitative results

Fig. 13 shows recall-precision results for protein-interaction extraction when tested on abstracts that have been manually tagged for protein names and Fig. 14 shows the results when tested on abstracts in which protein names were tagged using our best protein name extractor. We plotted a precision-recall curve for BWI by utilizing its extraction confidence, and for ELCS using the confidence levels of the rules which extract the interactions. Since RAPIER and human-written rules do not produce confidences, only a single recall-precision point is shown for each of them.

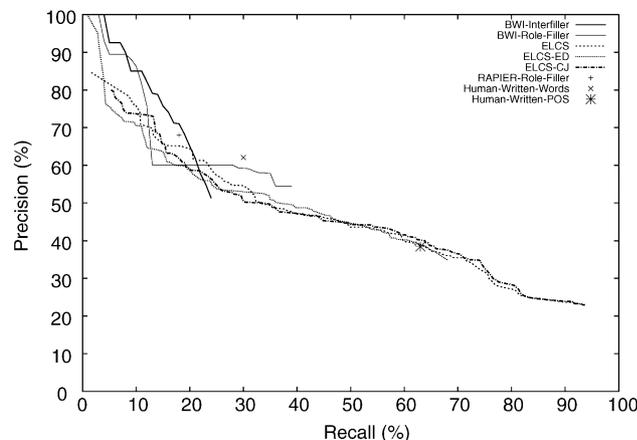


Figure 13 Precision-recall graphs for protein interaction extraction using correct protein names.

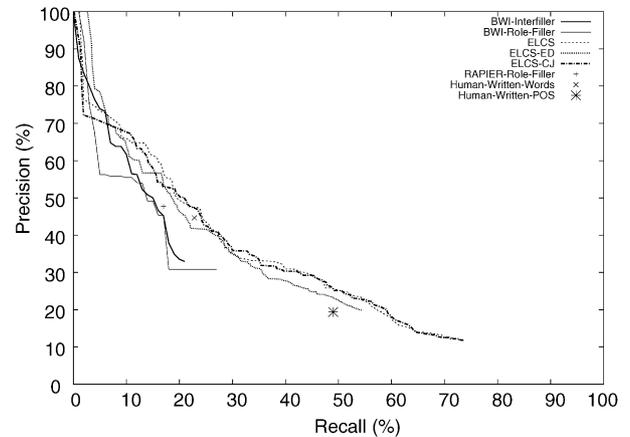


Figure 14 Precision-recall graphs for protein interaction extraction using extracted protein names.

4.2.3. Discussion of results

From Fig. 13 it can be seen that BWI gives varying degrees of high precision, but its recall is generally quite low. RAPIER also gives relatively high precision but low recall. ELCS tends to give higher recall with only a modest decrease in precision compared to BWI and RAPIER. When we use the protein names extracted from our protein name extractor instead of the correct protein names, not surprisingly the performance of all the systems degrade but they still offer reasonable ranges of precisions and recalls (Fig. 14).

These results demonstrate that machine learning systems can provide higher precisions than the human-written rules. In order to avoid over-loading human curators with too many false positives when extracting knowledge from large volumes of text, a general emphasis towards higher precision seems appropriate. The machine learning systems also offer a wide range of precision-recall trade-off which can be suitably utilized by a user depending on the needs of the application. The machine learning systems can also provide recalls higher than the best recall the human-written rules could provide. Hence, the machine learning systems out-perform the manually-written rules in several ways.

5. Conclusions and future research

After comparing a number of methods for extracting human protein names and interactions, we obtained the best performance for protein tagging with a maximum entropy learning method that exploits a generalized protein-name dictionary. For extracting protein interactions, we found that several methods for learning extraction rules out-perform the hand-written rules in providing higher precisions and in

offering a wider range along precision-recall trade-off. Token classification methods like k -NN, TBL, SVM, and MaxEnt are not directly applicable to extracting interactions; however, we plan to test HMMs on extracting interactions in the near future.

Clearly, the ability to extract human proteins and their interactions still needs significant improvement. We foresee improvement in three general areas: better training data, better learning methods, and better use of external knowledge.

Larger training sets are always beneficial to learning systems; however, manually tagging data is very time consuming. One alternative approach is to use existing knowledge to automatically produce *weakly labeled* training data [15]. Another approach is to use active learning to select only the best training examples for human labeling [53]. A third approach is to utilize a mixture of both labeled and unlabeled data during training [54].

Improved learning algorithms for information extraction continue to be developed. Recently, a number of methods for improving HMMs have been proposed, including linear interpolating HMMs [16], maximum-entropy Markov models [55], and conditional random fields [56].

Existing biological knowledge can also be used to improve extraction performance. Currently we have only exploited dictionaries of known protein names. Using learning to revise initial human-written extraction rules has also been shown to improve performance [18]. One can imagine many other sources of external knowledge: global statistical properties of abstracts, existing interaction or pathway data, prior expectations for finding protein names, and dictionaries of near-miss negative examples of protein names. Filtering proposed interacting proteins by comparing their gene-expression data or examining their co-occurrences in other abstracts or web pages could also prove useful.

In the future, it will also be interesting to develop IE systems for extracting other associations with genes. A few examples include extracting information about post-translational modifications of proteins, identifying genes that are specifically involved with diseases, identifying genes that are co-regulated, extracting protein-drug interactions, protein-metabolite interactions and information about the dynamics and dependencies of these processes.

Acknowledgements

We would like to thank members of the Marcotte lab for helping to tag Medline abstracts. We would also like to thank Kristie Seymore for making the IE-

tagging tool available. We would like to thank Lorraine Tanabe for kindly allowing us to use the ABGENE system. Many thanks to Kenichiro Fukuda for the KEX system, Mary Elaine Califf for the RAPIER system, Dayne Freitag and Nicholas Kushmerick for the BWI program, Thorsten Joachims for the SVM^{light} software, Eric Brill for his POS tagger, and to Ellen Riloff for the Sundance shallow parser. We also thank Christian Blaschke for the information about the hand-written rules. This work was supported in part by the National Science Foundation (IIS-0117308, ITR-0219061 and IIS-0325116), the Welch Foundation (F-1515), and the Texas Advanced Research Program.

References

- [1] Fukuda K, Tsunoda T, Tamura A, Takagi T. Information extraction: identifying protein names from biological papers. In: Proceedings of the 3rd Pacific Symposium on Biocomputing. Kapalua, HI; 1998. p. 707-18.
- [2] Humphreys K, Demetriou G, Geizauskas R. Two applications of information extraction to biological science journal articles: enzyme interactions and protein structure. In: Proceedings of the 5th Pacific Symposium on Biocomputing. Waikiki Beach, HI; 2000. p. 502-13.
- [3] Blaschke C, Valencia A. Can bibliographic pointers for known biological data be found automatically? protein interactions as a case study. *Compar Funct Genom* 2001;2:196-206.
- [4] Proux D, Rechenmann F, Julliard L. A pragmatic information extraction strategy for gathering data on genetic interactions. In: Proceedings of the 9th International Conference on Intelligent Systems for Molecular Biology (ISMB-2001). Copenhagen, Denmark; 2000. p. 279-85.
- [5] Rindflesch TC, Tanabe L, Weinstein JN, Hunter L. EDGAR: extraction of drugs, genes, and relations from the biomedical literature. In: Proceedings of the 5th Pacific Symposium on Biocomputing; 2000. p. 515-24.
- [6] Thomas J, Milward D, Ouzounis C, Pulman S, Carol M. Automatic extraction of protein interactions from scientific abstracts. In: Proceedings of the 5th Pacific Symposium on Biocomputing; 2000. p. 541-53.
- [7] Hahn U, Romacker M, Schulz S. Creating knowledge repositories from biomedical reports: the MEDSYNDIKATE text mining system. In: Proceedings of the 7th Pacific Symposium on Biocomputing; 2002. p. 338-49.
- [8] Pustejovsky J, Castano J, Zhang J, Kotecki M, Cochran B. Robust relational parsing over biomedical literature: extracting inhibit relations. In: Proceedings of the 7th Pacific Symposium on Biocomputing; 2002. p. 362-73.
- [9] Park J, Kim HS, Kim JJ. Bidirectional incremental parsing for automatic pathway identification with combinatory categorial grammar. In: Proceedings of the 6th Pacific Symposium on Biocomputing; 2001. p. 396-407.
- [10] Friedman C, Kra P, Yu H, Krauthammer M, Rzhetsky A. GENIES: a natural-language processing system for the extraction of molecular pathways from journal articles. *Bioinformatics* 2001;17(Suppl 1):S74-82.
- [11] Cowie J, Lehnert W. Information extraction. *Commun Assoc Comput Machinery* 1996;39(1):80-91.
- [12] Cardie C. Empirical methods in information extraction. *AI Mag* 1997;18(4):65-79.

- [13] Califf ME, editor. Papers from the AAAI-1999 Workshop on Machine Learning for Information Extraction. Orlando, FL: AAAI Press; 1999.
- [14] Proux D, Rechenmann F, Julliard L, Pillet V, Jacq B. Detecting gene symbols and names in biological texts: a first step toward pertinent information extraction. *Genome Inform* 1998;9:72–80 (GIW'98).
- [15] Craven M, Kumlien J. Constructing biological knowledge bases by extracting information from text sources. In: Proceedings of the 7th International Conference on Intelligent Systems for Molecular Biology (ISMB-1999); 1999. p. 77–86.
- [16] Collier N, No C, Tsujii J. Extracting the names of genes and gene products with a hidden Markov model. In: Proceedings of the 18th International Conference on Computational Linguistics, Saarbrücken; 2000. p. 201–7.
- [17] Ray S, Craven M. Representing sentence structure in hidden Markov models for information extraction. In: Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI-2001); 2001. p. 1273–9.
- [18] Eliassi-Rad T, Shavlik J. A theory-refinement approach to information extraction. In: Proceedings of the 18th International Conference on Machine Learning (ICML-2001); 2001.
- [19] Leonard JE, Colombe JB, Levy JL. Finding relevant references to genes and proteins in medline using a Bayesian approach. *Bioinformatics* 2002;18(11):1515–22.
- [20] Raychaudhuri S, Chang JT, Sutphin PD, Altman RB. Associating genes with gene ontology codes using a maximum entropy analysis of biomedical literature. *Genome Res* 2002;12:203–14.
- [21] Perez-Iratxeta C, Bork P, Andrade MA. Association of genes to genetically inherited diseases using data mining. *Nat Genet* 2002;31(3):316–9.
- [22] Venter, JC, et al. The sequence of the human genome. *Science* 16 February 2001;291(5507):1304–51.
- [23] Lander, ES, et al. Initial sequencing and analysis of the human genome. *Nature* 15 February 2001;409(6822):860–921.
- [24] Califf ME, Mooney RJ. Relational learning of pattern-match rules for information extraction. In: Proceedings of the 16th National Conference on Artificial Intelligence (AAAI-99); 1999. p. 328–34.
- [25] Freitag D, Kushmerick N. Boosted wrapper induction. In: Proceedings of the 17th National Conference on Artificial Intelligence (AAAI-2000). Austin, TX: AAAI Press/The MIT Press; 2000. p. 577–83.
- [26] Duda RO, Hart PE. Pattern classification and scene analysis. New York: Wiley; 1973.
- [27] Brill E. Transformation-based error-driven learning and natural language processing: a case study in part-of-speech tagging. *Computat Linguist* 1995;21(4):543–65.
- [28] Vapnik VN. The nature of statistical learning theory. Berlin: Springer-Verlag; 1995.
- [29] Berger AL, Della Pietra SA, Della Pietra VJ. A maximum entropy approach to natural language processing. *Computat Linguist* 1996;22(1):39–71.
- [30] Xenarios I, Fernandez E, Salwinski L, Duan XJ, Thompson MJ, Marcotte EM, Eisenberg D. DIP: The database of interacting proteins: 2001 update. *Nucleic Acids Res* 2001;29(1):239–41.
- [31] Bikel DM, Schwartz R, Weischedel RM. An algorithm that learns what's in a name. *Machine Learning* 1999;34:211–32.
- [32] Roth D, van den Bosch A, editors. Proceedings of the 6th Conference on Natural Language Learning (CoNLL-2002). Association for Computational Linguistics. Taipei, Taiwan; 2002.
- [33] Kushmerick N, Weld DS, Doorenbos RB. Wrapper induction for information extraction. In: Proceedings of the 15th International Joint Conference on Artificial Intelligence (IJCAI-97); 1997. p. 729–35.
- [34] Freund Y, Schapire RE. Experiments with a new boosting algorithm. In: Saitta L, editor. Proceedings of the 13th International Conference on Machine Learning (ICML-96), 1996.
- [35] Vapnik VN. Statistical learning theory. New York: John Wiley & Sons; 1998.
- [36] Joachims T. Text categorization with support vector machines: learning with many relevant features. In: Proceedings of the 10th European Conference on Machine Learning. Berlin: Springer-Verlag; 1998. p. 137–42.
- [37] Ramshaw LA, Marcus MP. Text chunking using transformation-based learning. In: Proceedings of the 3rd Workshop on Very Large Corpora; 1995.
- [38] Kudoh T, Matsumoto Y. Use of support vector learning for chunk identification. In: Proceedings of CoNLL-2000 and LLL-2000. p. 142–4.
- [39] Kubat M, Holte RC, Matwin S. Machine learning for the detection of oil spills in satellite radar images. *Machine Learning* 1998;30(2–3):195–215.
- [40] Collins M. Ranking algorithms for named-entity extraction: boosting and the voted perceptron. In: Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL-02); 2002.
- [41] Rabiner LR. A tutorial on hidden Markov models and selected applications in speech recognition. *Proc IEEE* 1989;77(2): 257–86.
- [42] Darroch J, Ratchliff D. Generalized iterative scaling for log-linear models. *Ann Math Stat* 1972;43(5):1470–80.
- [43] Tanabe L, Wilbur WJ. Tagging gene and protein names in biomedical text. *Bioinformatics* 2002;18(8):1124–32.
- [44] Marcus M, Santorini B, Marcinkiewicz MA. Building a large annotated corpus of English: the Penn treebank. *Computat Linguist* 1993;19(2):313–30.
- [45] National Institute of Standards and Technology, ACE—automatic content extraction (<http://www.nist.gov/speech/tests/ace/>)
- [46] Marcotte E, Xenarios I, Eisenberg D. Mining literature for protein–protein interactions. *Bioinformatics* April 2001;17(4):359–63.
- [47] Blaschke C, Valencia A. The frame-based module of the Suiseki information extraction system. *IEEE Intell Sys* 2002;17:14–20.
- [48] Charras C, Lecroq T. Sequence comparison, Laboratoire d'Informatique de Rouen et Atelier Biologie Informatique Statistique Socio-Linguistique, Université de Rouen, France; 1998 (<http://www-igm.univ-mlv.fr/lecroq/seqcomp>)
- [49] Gusfield D. Algorithms on strings, trees and sequences. New York: Cambridge University Press; 1997.
- [50] Levenshtein VI. Binary codes capable of correcting insertions and reversals. *Sov Phys Doklady* 1966;10(8):707–10.
- [51] Cestnik B. Estimating probabilities: a crucial task in machine learning. In: Proceedings of the 9th European Conference on Artificial Intelligence; 1990. p. 147–9.
- [52] Pearl J. Probabilistic reasoning in intelligent systems: networks of plausible inference. San Mateo, CA: Morgan Kaufmann; 1988.
- [53] Thompson CA, Califf ME, Mooney RJ. Active learning for natural language parsing and information extraction. In: Proceedings of the 16th International Conference on Machine Learning (ICML-99); 1999. p. 406–14.
- [54] M. Collins, Y. Singer, Unsupervised models for named entity classification. In: Proceedings of the Conference on Empiri-

- cal Methods in Natural Language Processing and Very Large Corpora (EMNLP/VLC-99). University of Maryland, College Park, MD; 1999.
- [55] McCallum A, Freitag D, Pereira F. Maximum entropy Markov models for information extraction and segmentation. In: Proceedings of the 17th International Conference on Machine Learning (ICML-2000); 2000.
- [56] J. Lafferty, A. McCallum, F. Pereira, Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In: ICML-2001. Williamstown, MA; 2001.