

Improving the Performance of Dictionary-based Approaches in Protein Name Recognition

Yoshimasa Tsuruoka^{††} and Jun'ichi Tsujii^{†‡}

[†]*Department of Computer Science, University of Tokyo*

Hongo 7-3-1, Bunkyo-ku, Tokyo 113-0033 Japan

[‡]*CREST, JST (Japan Science and Technology Agency)*

Honcho 4-1-8, Kawaguchi-shi, Saitama 332-0012 Japan

`{tsuruoka,tsujii}@is.s.u-tokyo.ac.jp`

Phone : +81-3-5803-1697

FAX : +81-3-5802-8872

Abstract

Dictionary-based protein name recognition is often a first step in extracting information from biomedical documents because it can provide ID information on recognized terms. However, dictionary-based approaches present two fundamental difficulties: (1) false recognition mainly caused by short names; (2) low recall due to spelling variations. In this paper, we tackle the former problem using machine learning to filter out false positives and present two alternative methods for alleviating the latter problem with spelling variations. The first is achieved by using approximate string searching, and the second by expanding the dictionary with a probabilistic variant generator, which we propose in this paper. Experimental results using the GENIA corpus revealed that filtering using a naive Bayes classifier greatly improved precision with only a slight loss of recall, resulting in 10.8% improvement in F-measure, and dictionary expansion with the variant generator gave further 1.6% improvement and achieved an F-measure of 66.6%.

Key words: protein name recognition, naive Bayes classifier, approximate string search, spelling variant generator

1 Introduction

The rapid increase in machine readable biomedical texts (e.g. MEDLINE) makes automatic information extraction from these texts much more attractive. One of the most important tasks today is extracting information on protein-protein interactions from MEDLINE abstracts (1; 2; 3).

To be able to extract information about proteins from a text, one has to first recognize their names in it. This kind of problem has been extensively studied in the field of natural language processing as named-entity recognition tasks. The most popular approach is to train the recognizer on an annotated corpus by using a machine learning algorithm, such as Hidden Markov Models, support vector machines (SVMs) (4), and maximum-entropy models (5). The task of the classifier in the machine learning framework is to determine the text regions corresponding to protein names.

Ohta et al. provided the GENIA corpus (6), an annotated corpus of MEDLINE abstracts, which could be used as a gold-standard for evaluating and training named-entity recognition algorithms. The corpus has fostered research on machine learning techniques for recognizing biological entities in texts (7; 8; 9).

However, the main drawback of these machine learning approaches is that they do not provide ID information on recognized terms. For the purpose of extracting information about proteins, ID information on recognized proteins, such as GenBank ¹ ID or SwissProt ² ID, is indispensable to integrate extracted

¹ GenBank is one of the largest genetic sequence databases.

² The Swiss-Prot is an annotated protein sequence database.

information with relevant data from other information sources.

Dictionary-based approaches, on the other hand, intrinsically provide ID information because they recognize a term by searching the one that is most similar (or identical) in the dictionary to the target region. This advantage makes dictionary-based approaches particularly useful as the first step in practical information extraction from biomedical documents (3).

Dictionary-based approaches, however, present two fundamental difficulties. The first is a large number of false positives mainly caused by short names, which significantly degrades overall precision. Although this problem can be avoided by excluding short names from the dictionary, such a solution makes it impossible to recognize short protein names. We tackle this problem by incorporating a machine learning technique to filter out the false positives.

The other problem in dictionary-based approaches derives from the fact that biomedical terms have many spelling variations. For example, the protein name “NF-Kappa B” has many spelling variants such as “NF Kappa B”, “NF kappa B”, “NF kappaB” and “NFkappaB”. Exact matching techniques regard these terms as completely different terms, which results in failure to find these protein names written in various forms.

We present two alternative solutions to the problem of spelling variation in this paper. The first is using approximate string searching techniques where the surface-level similarity of strings is considered. The second is expanding the dictionary in advance with a *probabilistic variant generator*, which we propose in this paper. We present experimental results on the GENIA corpus to demonstrate their effectiveness,

This paper is organized as follows. Section 2 overviews our method of recognizing protein names. Section 3 explains the approximate string searching algorithm used to alleviate the problem of spelling variations. As an alternative solution to the problem, Section 4 describes the probabilistic variant generator that is used to expand the dictionary. Section 5 describes how false recognition is filtered out with a machine learning method. Section 6 presents experimental results obtained using the GENIA corpus. Some related work is described in Section 7. Finally, Section 8 has some concluding remarks.

2 Method Overview

Our method of recognizing protein names involves the following two phases.

- Candidate recognition phase

The task of this phase is to find protein name candidates appearing in the text using a dictionary. We propose two alternative solutions to the problem with spelling variations. The first is to use an approximate string searching algorithm instead of exact matching algorithms, which is presented in Section 3. The second is to expand the dictionary in advance with the variant generator, which is presented in Section 4.

- Filtering phase

One of the most serious problems with dictionary-based recognition is the large number of false positives mainly caused by short entries in the dictionary. Our solution to this problem is to check whether each candidate is really a protein name or not. In other words, each protein name candidate is classified into “accepted” or “rejected” with a machine learning algorithm. The classifier uses the context of the term and the term itself as

	G R - 2				
	0	1	2	3	4
E	1	1	2	3	4
G	2	1	2	3	4
R	3	2	1	2	3
-	4	3	2	1	2
1	5	4	3	2	2

Fig. 1. Dynamic programming matrix.

the features for classification. Only “accepted” candidates are recognized as protein names in the final output. Section 5 describes details in the classification algorithm used in this phase.

In the following sections, we describe details on the methods used in these phases.

3 Candidate Recognition by Approximate String Searching

One way to deal with the problem of spelling variations is to use a kind of “elastic” matching algorithm, by which a recognition system scans a text to find a similar term to (if any) a protein name in the dictionary. We need a similarity measure for this task. The most popular measure of similarity between two strings is the *edit distance*, which is the minimum number of operations on individual characters (e.g. substitutions, insertions, and deletions) required to transform one string of symbols into another. For example, the edit distance between “EGR-1” and “GR-2” is two, because one substitution (1 for 2) and one deletion (E) are required.

To calculate the edit distance between two strings, we can use a dynamic programming technique (10). Figure 1 illustrates an example ³. In matrix $C_{0..|x|,0..|y|}$, each cell $C_{i,j}$ keeps the minimum number of operations needed to match $x_{1..i}$ to $y_{1..j}$ and can be computed as a simple function of the surrounding cells:

$$C_{i,0} = i, \tag{1}$$

$$C_{0,j} = j, \tag{2}$$

$$C_{i,j} = \begin{cases} (x_i = y_j) & \text{then } C_{i-1,j-1} \\ \text{else } 1 + \min(C_{i-1,j}, C_{i,j-1}, C_{i-1,j-1}). \end{cases} \tag{3}$$

The calculation can be done either in row-wise left-to-right traversal or in column-wise top-to-bottom traversal.

There are some algorithms that run faster than the dynamic programming method in computing uniform-cost edit distance, where the weight of each edit operation is constant within the same type (11). However, what we expect is that the distance between “EGR-1” and “EGR 1” will be smaller than that between “EGR-1” and “FGR-1”, while their uniform-cost edit distances are equal.

The dynamic programming based method is flexible enough to allow us to define arbitrary costs for individual operations depending on the letter being operated on. For example, we can make the cost of a substitution between a space and a hyphen much lower than that of a substitution between ‘E’ and ‘F’. Therefore, we use the dynamic-programming-based method for our task.

Table 1 lists the cost functions we used in our experiments. As we can see, both insertion and deletion costs are 100 except for spaces and hyphens. Substitu-

³ For clarity of presentation, all costs have been assumed to be 1.

Table 1

Cost function.

Operation	Letter	Cost
Insertion	<i>a space or a hyphen</i>	10
	other letters	100
Deletion	<i>a space or a hyphen</i>	10
	other letters	100
Substitution	a numeral for a numeral	10
	<i>a space for a hyphen</i>	10
	<i>a hyphen for a space</i>	10
	a capital letter for the corresponding small letter	10
	a small letter for the corresponding capital letter	10
	other letters	50

tion costs for similar letters are 10. Substitution costs for the other different letters are 50. Since these costs were heuristically determined by just observing a number of protein names, it is likely that we could achieve better performance by employing a systematic method of determining the cost function (12).

3.1 String Searching

What we described in the previous section is a method of calculating the similarity between two strings. However, what we need when trying to find proteins where is approximate string searching in which the recognizer scans a text to find a similar term to (if any) a term in the dictionary. The dynamic-programming-based method can be easily extended for approximate string searching.

The method is illustrated in Figure 2. In this case, the protein name to be

	e n c o d e d b y E G R 1 i n c l u d e																											
E	0	1	2	3	4	5	6	7	0	1	2	0	1	2	3	0	1	0	1	2	3	4	5	6	7			
G	1	1	2	3	4	5	6	7	1	1	2	1	0	1	2	1	1	1	1	2	3	4	5	6	7			
R	2	2	2	3	4	5	6	7	2	2	2	2	1	0	1	2	2	2	2	2	3	4	5	6	7			
-	3	3	3	3	4	5	6	7	3	3	3	3	2	1	0	1	2	3	3	3	3	3	4	5	6	7		
1	4	4	4	4	4	5	6	7	4	4	4	4	3	2	1	1	2	3	4	4	4	4	4	5	6	7		
	5	5	5	5	5	5	6	7	5	5	5	5	4	3	2	2	1	2	3	4	5	5	5	6	7			

Fig. 2. String searching using dynamic programming matrix.

matched is “EGR-1” and the text to be scanned is “encoded by EGR include”. String searching can be done by just setting the elements corresponding to the separators (e.g. space) in the first row to zero. After filling the whole matrix, one can find that “EGR-1” can be matched to this text at the place of “EGR 1” with cost 1 by searching for the lowest value in the bottom row and then backtracing to the top row along the lowest-cost path.

To take into account the length of a term, we use a normalized cost, which is calculated by dividing the cost by the length of the term:

$$(\text{normalized cost}) = \frac{(\text{cost}) + \alpha}{(\text{length of the term})}, \quad (4)$$

where α is a constant value ⁴. When the costs for two terms are equal, the longer one is preferred due to this constant. In the case of Figure 2, the normalized cost of the term is $(1 + 0.4)/4 = 0.35$.

To recognize a protein name in a given text, we do the above calculation for every term contained in the dictionary and select the term that has the

⁴ α was heuristically set to 0.4 in our experiments. It would be possible to tune the value by conducting cross-validation on the training data with additional computational costs.

lowest normalized cost. If the cost is lower than the predefined threshold, the corresponding range in the text is recognized as a protein name candidate.

3.2 Implementation Issues in String Searching

A naive way of string searching using a dictionary is to go through the procedure described in the previous section one by one for every term in the dictionary. However, since a protein name dictionary is usually large ($\sim 10^5$), this naive way requires too much computational cost to deal with a large number of documents.

Navarro et al. (13) presented a way of reducing redundant calculations by constructing a trie of the dictionary. The trie is used as a device to avoid repeating the computation of the cost against the same prefix of many patterns. Suppose that we have just calculated the cost of the term “EGR-1” and we next have to calculate the cost of the term “EGR-2”; it is clear that we do not have to re-calculate the first four rows in the matrix (see Figure 2). They also pointed out that it is possible to determine, prior to reaching the bottom of the matrix, that the current term cannot produce any relevant match: if all the values of the current row are larger than the threshold, then a match cannot occur since we can only increase the cost or at best keep it the same.

We followed their method: we first constructed a trie and fill the matrix one by one for each term, avoiding redundant calculations. The computational cost for approximate string searching was very large even with these devices for efficient computation. In our implementation, it took more than an hour to process 2,000 MEDLINE abstracts on a 1.13 GHz pentium III server.

4 Expanding Dictionary with Probabilistic Variant Generator

An alternative way to alleviate the problem of spelling variations is to expand each entry in the dictionary in advance. For example, if we have the entry “EGR-1” in the dictionary, we expand this entry to the two entries “EGR-1” and “EGR 1”. With the expanded dictionary, we can find protein names written in varied forms simply by using exact-matching algorithms.

For this purpose, we propose an algorithm that can generate only “likely” spelling variants. Our method not only generates spelling variants but also gives each variant a *generation probability* that represents the plausibility of the variant. Therefore, one does not need to receive a prohibitive number of unnecessary variants by setting an appropriate threshold for generation probability.

4.1 Probabilistic Variant Generator

4.1.1 Generation Probability

The *generation probability* of a variant is defined as the probability that the variant can be generated through a sequence of operations. Each operation has an *operation probability* that represents how likely it is that it will occur. Assuming independence among operations, the generation probability of a variant can be formalized in a recursive manner,

$$P_X = P_Y \times P_{op}, \tag{5}$$

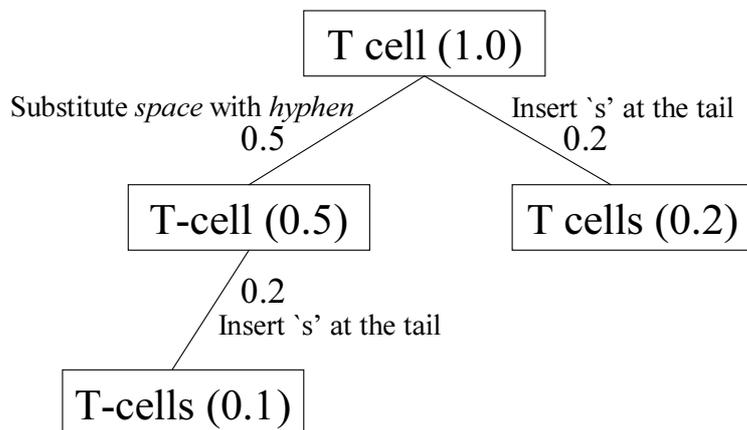


Fig. 3. Probabilistic variant generation. Numerals inside parentheses are generation probabilities, and those along the edges are operation probabilities.

where P_X is the generation probability of variant X , P_Y is the generation probability of variant Y from which variant X is generated, and P_{op} is the probability of the operation by which Y is transformed into X .

Figure 3 outlines an example of the generation process, which can be represented as a tree. Each node represents a generated variant and its probability. Each edge represents an operation and its probability. The root node corresponds to the input term and the generation probability of the root node is 1 by definition. We can obtain the variants of an input term in order of their generation probabilities by growing a tree in a best-first manner.

4.1.2 Operation Probability

To calculate the generation probabilities in our formalization, we need the probability for each operation.

We used three types of operations for the generation mechanism:

- Substitution

Replace a character with another character.

- Deletion

Delete a character.

- Insertion

Insert a character.

These types of operations are motivated by the ones used in approximate string matching. We consider character-level contexts in which an operation occurs, and the following seven types of contexts are used in this paper. They differ in relative position to the target and in how much the context is specified:

- the target letter and the preceding two letters.
- the target letter and the preceding letter.
- the target letter and the following letter.
- the target letter and the following two letters.
- the target letter, the preceding letter, and the following letter.
- the target letter, the preceding two letters, and the following two letters.
- the target letter only.

For an operation for a substitution or a deletion, the target indicates a letter in the string. For an operation for an insertion, the target indicates a gap between two letters. For example, if the original string is “c-Rel” and the variant is “c-rel”, the operation is a substitution of ‘R’ with ‘r’. The operation rules obtained from this example are listed in Table 2. They correspond to the seven types of aforementioned context. The first rule indicates that if the letter ‘R’ is preceded by the string “c-”, then one can replace the letter with ‘r’.

The next step is estimating the probability of each rule. The probability should

Table 2

Example of operation rules.

Left Context	Target	Right Context	Operation
c-	R	*	Replace the target with 'r'
-	R	*	Replace the target with 'r'
*	R	e	Replace the target with 'r'
*	R	el	Replace the target with 'r'
-	R	e	Replace the target with 'r'
c-	R	el	Replace the target with 'r'
*	R	*	Replace the target with 'r'

Asterisks represent wild cards.

represent how likely the operation is to occur in a given context. We estimate the operation probabilities from a large number of variant pairs with the following equation:

$$P_{op} = P(\text{operation}|\text{context}) \quad (6)$$

$$\approx \frac{f(\text{context}, \text{operation}) + 1}{f(\text{context}) + 2}, \quad (7)$$

where $f(\text{context})$ is the frequency of the occurrence of the context, and $f(\text{context}, \text{operation})$ is the frequency of the simultaneous occurrence of the context and operation in the set of variant pairs. We adopt Laplace smoothing (adding 1 to the numerator and 2 to the denominator).

We acquired the actual samples of variant pairs for probability estimation from the UMLS Metathesaurus (14), which provides a huge number of biomedical terms and their concept IDs. We first obtained sets of variants by collecting protein names with the same concept ID. We then selected from each set the pairs whose edit distance is one, and used them for probability estimation.

For example, we could obtain the following variant pairs from the set shown in Table 3.

Table 3

A part of UMLS Metathesaurus.

Concept ID	Protein name
:	:
C0079930	Oncogene Protein gp140(v-fms)
C0079930	Oncogene protein GP140, V-FMS
C0079930	fms Oncogene Product gp140
C0079930	fms Oncogene Protein gp140
C0079930	gp140(v-fms)
C0079930	gp140 v fms
C0079930	gp140 v-fms
C0079930	v-fms, gp140
C0079930	v-fms Protein
C0079930	V-FMS protein
C0079930	v fms Protein
C0079930	Oncogene protein V-FMS
C0079930	GP140 V-FMS protein
:	:

{“gp140 v fms”, “gp140 v-fms”}

{“v-fms Protein, “v fms Protein”}

4.1.3 Generation Algorithm

Once the rules and their probabilities are learned, we can generate variants from an input term using those rules.

The whole algorithm for variant generation is given below. Note that V represents the set of generated terms.

- (1) Initialization

Add the input term to V .

(2) Selection

Select the term and the operation to be applied to it so that the algorithm will generate a new term which has the highest possible probability.

(3) Generation

Generate a new term using the term and the operation selected in Step 2. Then, add the generated term to V .

(4) Repeat

Go back to Step 2 until the termination condition is satisfied.

In the generation step, the system applies the rule whose context matches any part of the string. If multiple rules can be applied, the rule that has the highest operation probability is used.

Because this algorithm generates variants in the order of their generation probabilities, the termination condition can be where that the generation probability for the generated variant is below the predefined threshold or that the number of generated variants exceeds the predefined threshold.

5 Filtering Candidates by Naive Bayes Classifier

In the filtering phase, we use a classifier trained on an annotated corpus to suppress false recognition. The objective of this phase is to improve precision without the loss of recall.

We conduct binary classification (“accept” or “reject”) on each candidate. The candidates that are classified into “rejected” are filtered out. In other words, only the candidates that are classified into “accepted” are recognized

as protein names in the final output.

In this paper, we use a naive Bayes classifier for this task of classification. The naive Bayes model is simple but effective and has been used in numerous applications of information processing including image recognition, natural language processing, and information retrieval (15; 16; 17; 18). The model assumes conditional independence among features, and the assumption makes it possible to estimate the model parameters from limited amount of training data.

There are some implementation variants with the naive Bayes classifier depending on their event models (19). In this paper, we adopted the multi-variate Bernoulli event model, in which all features are binary.

5.1 Features

The features of the target must be represented in the form of a vector. for the input of the classifier. We use a binary feature vector which contains only the values of 0 or 1 for each element.

In this paper, we use the local context surrounding a candidate term and the words contained in the term as the features. We call the former *contextual features* and the latter *term features*.

The features used in our experiments are given below.

- Contextual Features

W_{-1} : the preceding word.

W_{+1} : the following word.

- Term Features

W_{begin} : the first word of the term.

W_{end} : the last word of the term.

W_{middle} : the other words of the term without positional information (bag-of-words).

Suppose the candidate term is “putative zinc finger protein” and the sentence is

... encoding a putative zinc finger protein was found to derepress beta-galactosidase ...

We obtain the following active features for this example.

$\{W_{-1}$ a $\}$, $\{W_{+1}$ was $\}$, $\{W_{begin}$ putative $\}$, $\{W_{end}$ protein $\}$, $\{W_{middle}$ zinc $\}$,
 $\{W_{middle}$ finger $\}$.

5.2 Training

The training of the classifier is done on an annotated corpus. We first scan the corpus for protein name candidates with the dictionary-matching method described in Section 3 or 4. If a recognized candidate is annotated as a protein name, this candidate and its context are used as a positive (“accepted”) example for training. Otherwise, it is used as a negative (“rejected”) example.

6 Experiment

6.1 Corpus and Dictionary

We conducted experiments on protein name recognition using the GENIA corpus version 3.02 (6), which contains 2,000 abstracts extracted from the MEDLINE database. These abstracts were selected from the search results with the MeSH terms *Human*, *Blood Cells*, and *Transcription Factors*.

The biological entities in the corpus are annotated according to the GENIA ontology. Although the corpus has many categories such as protein, DNA, RNA, cell line and tissue, we only used the protein category. When a term was recursively annotated, only the outermost (longest) annotation was used.

We conducted ten-fold cross-validation for evaluating the methods. Each set of 200 abstracts was used as the test data, and the remaining 1,800 abstracts were used as the training data. The results were averaged over the ten runs.

The protein name dictionary was constructed from the training data by collecting all the terms that were annotated as proteins ⁵. The average number of terms contained in the dictionary was 8,055.

Each recognition was counted as correct if both boundaries of the recognized term exactly matched the boundaries of an annotation in the corpus.

⁵ We used the training data for constructing the protein name dictionary so that the recognition performance could be comparable with those of machine learning based approaches. However, there is a chance that we have lower recognition performance if we use an external resource for constructing the dictionary.

6.2 Improving Precision by Filtering

We first conducted experiments to evaluate how much precision had been improved by the filtering process. In the candidate-recognition phase, the longest matching algorithm was used for candidate recognition.

The results are listed in Table 4. The F-measure is defined as the harmonic mean of precision and recall:

$$F = \frac{2 \times \textit{precision} \times \textit{recall}}{\textit{precision} + \textit{recall}} \quad (8)$$

The first row shows performance achieved without filtering where all candidates identified in the candidate-recognition phase are regarded as protein names. The other rows shows performance achieved with filtering using the naive Bayes classifier. In this case, only the candidates that were classified into “accepted” were regarded as protein names. Notice that the filtering significantly improved the precision (from 46.5% to 71.2%) with only a slight loss of recall. The F-measure was also greatly improved (from 54.3% to 65.1%).

6.2.1 Efficacy of Contextual Features

The advantage of using a machine learning technique is that we can exploit the context of a candidate. To evaluate the efficacy of contexts, we conducted experiments using different feature sets.

The results in Table 4 indicate that candidate terms themselves provide strong cues for classification. However, the fact that the best performance was achieved when both feature sets were used suggests that the context of a candidate term conveys useful information about its the semantic class.

Table 4

Precision improvement by filtering.

	Feature Sets	Precision	Recall	F-measure
Without filtering	N/A	46.5%	65.4%	54.3%
With filtering	Contextual features	60.1%	58.0%	59.0%
	Term features	68.2%	59.8%	63.7%
	All features	71.2%	60.1%	65.1%

Table 5

Effectiveness of approximate string search.

	Without Filtering			With Filtering		
Threshold	Precision	Recall	F-measure	Precision	Recall	F-measure
4	45.8%	67.4%	54.5%	70.6%	61.8%	65.9%
8	42.4%	68.4%	52.3%	69.6%	63.0%	66.1%
12	37.4%	69.3%	48.5%	68.3%	63.9%	66.0%
16	30.6%	70.0%	42.5%	66.7%	64.7%	65.6%
20	21.0%	71.4%	32.4%	63.5%	66.1%	64.7%

6.3 Improving Recall by Approximate String Search

We conducted experiments to evaluate how much we could further improve recognition by using the approximate string-searching method described in Section 3. Table 5 lists the results. The leftmost columns show the thresholds of normalized costs for approximate string searching. As the threshold increased, the precision degraded while the recall improved. The best F-measure was 66.1%, which is better than that of exact matching by 1.0% (see Table 4). It should be noted that approximate string-searching without filtering did not improve F-measure: the decline of precision overwhelmed the improvement of recall.

6.4 *Expanding Dictionary by Variant Generation*

6.4.1 *Variant Generator*

We used the UMLS Metathesaurus, version 2003AA to learn the operation rules. Terms having the semantic type of “Amino Acid, Peptide, or Protein” were used for learning, and they provide 36,112 pairs of spelling variants. Table 6 lists some of the rules obtained their probabilities. The asterisks in the table represent wild cards, meaning that one can ignore the context of that position. Notice that there are many rules for replacing spaces and hyphens. This suggests that spaces and hyphens are often used interchangeably in protein names.

The variants of some biomedical terms generated by our algorithm are listed in Tables 7 to 10. The first two input terms “NF-kappa B” and “transcription factor” are the two most frequent protein names in the corpus.

The generated variants of “transcription factor” listed in Table 8 are interesting. The first letters of “transcription” and “factor” were capitalized in the second and third variant respectively. This reflects the fact that the first letter of a word is often capitalized in biomedical terms. Notice that the plural form of “factor” is generated in the first variant.

The variants for the input term “tumor necrosis factor” are listed in Table 9. It should be noted that transformation to the British spelling variation for “tumor” appears in the seventh variant.

The variants for the input term “T cell factor 1” are listed in Table 10. Note that the variant where a hyphen is inserted between the ‘T’ and “cell” is

Table 6

Operation rules and their probabilities.

Operation Probability	Left Context	Target	Right Context	Operation
0.971	*	^	*	delete the target
0.958	*	o	ea	delete the target
0.958	rh		ea	insert 'o'
0.952	e	<i>hyphen</i>	R	replace the target with <i>space</i>
0.950	or	<i>space</i>	3,	replace the target with <i>hyphen</i>
0.950	or	<i>hyphen</i>	3,	replace the target with <i>space</i>
0.947	TH		(1	insert <i>space</i>
0.947	or	s	_a	delete the target
0.945	*	<i>space</i>	tR	replace the target with <i>hyphen</i>
0.938	_T	<i>space</i>	Ce	replace the target with <i>hyphen</i>
0.938	L	<i>space</i>	I	replace the target with <i>hyphen</i>
0.938	in	<i>space</i>	bi	replace the target with <i>hyphen</i>
0.938	ne	<i>space</i>	tR	replace the target with <i>hyphen</i>
0.938	3	<i>space</i>	*	replace the target with <i>hyphen</i>
0.938	r	<i>hyphen</i>	3	replace the target with <i>space</i>
0.938	E	<i>space</i>	1	replace the target with <i>hyphen</i>
0.938	V	<i>space</i>	I	replace the target with <i>hyphen</i>
0.938	ne	s	_R	delete the target
0.938	*	s	_2	delete the target
0.929	<i>start of term</i>	l	o	replace the target with 'L'
0.929	NA	<i>space</i>	DE	replace the target with <i>hyphen</i>
0.929	NA	<i>hyphen</i>	DE	replace the target with <i>space</i>
0.929	in	<i>hyphen</i>	A	replace the target with <i>space</i>
0.929	rg	<i>hyphen</i>	*	replace the target with <i>space</i>
0.923	k	a	e	delete the target
0.923	x	<i>hyphen</i>	1	delete the target
:	:	:	:	:

Asterisks represent wild cards.

Table 7

Generated variants for “NF-kappa B”.

Generation Probability	Generated String
1.000	NF-kappa B
0.466	NF kappa B
0.317	NF-kappa-B
0.286	NF-Kappa B
0.233	NFkappa B
0.211	NP-kappa B
0.199	NP kappa B
0.190	NF Kappa B
0.150	NF-kappaB
0.148	NF kappa-B
0.090	NF-Kappa-B
0.081	NP Kappa B
:	:

ranked at the top. The eighth variant “T cell factor I” is also interesting, where the numeral ‘1’ is replaced with the letter ‘I’.

6.4.2 Dictionary Expansion

We conducted experiments on dictionary expansion using the variant generator. Expansions were done on the terms whose lengths were equal to or longer than five characters. The maximum number of variants generated for each term was limited to 100 ⁶.

Table 11 shows the effectiveness of dictionary expansion. The leftmost columns show the threshold for generation probability to expand the dictionary. The

⁶ For equi-probable variants, the program continues to generate even when the number of variants exceeds the limit.

Table 8

Generated variants for “transcription factor”.

Generation Probability	Generated String
1.000	transcription factor
0.571	transcription factors
0.356	Transcription factor
0.219	transcription Factor
0.206	trancription factor
0.203	Transcription factors
0.137	transcription-factor
0.125	transcription Factors
0.117	trancription factors
0.107	transcription factorss
0.078	transcription-factors
0.073	Trancription factor
:	:

Table 9

Generated variants for “Tumor necrosis factor”.

Generation Probability	Generated String
1.000	Tumor necrosis factor
0.571	Tumor necrosis factors
0.218	Tumor necrosi factor
0.188	Tumors necrosis factor
0.176	tumor necrosis factor
0.139	Tumor-necrosis factor
0.137	Tumor necrosis-factor
0.125	Tumour necrosis factor
0.124	Tumor necrosis Factor
0.124	Tumor necrosi factors
0.107	Tumors necrosis factors
:	:

Table 10

Generated variants for “T cell factor 1”.

Generation Probability	Generated String
1.000	T cell factor 1
0.604	T-cell factor 1
0.498	T cell factor-1
0.301	T-cell factor-1
0.196	T cell factors 1
0.139	T cell factor1
0.137	T cell-factor 1
0.135	t cell factor 1
0.129	T cell factor I
0.124	T cell Factor 1
0.118	T-cell factors 1
:	:

Table 11

Effectiveness of Dictionary Expansion

Threshold	Without Filtering			With Filtering		
	Precision	Recall	F-measure	Precision	Recall	F-measure
2^{-1}	46.4%	66.4%	54.6%	71.2%	61.0%	65.7%
2^{-2}	46.5%	67.1%	54.9%	71.6%	61.7%	66.2%
2^{-3}	46.5%	67.4%	55.0%	71.6%	61.7%	66.3%
2^{-4}	46.4%	67.7%	55.0%	71.8%	62.1%	66.5%
2^{-5}	46.0%	68.0%	54.8%	71.7%	62.3%	66.6%
2^{-6}	46.0%	68.1%	54.8%	71.7%	62.3%	66.6%

recall improved as the threshold decreased. The best F-measure was 66.6%, which is 1.6% higher than that with the original dictionary (see Table 4).

7 Related Work

Kazama et al. (9) reported an F-measure of 56.5% on the GENIA corpus version 1.1 using SVMs. Collier et al. (20) reported an F-measure of 75.9% on 100 MEDLINE abstracts using a Hidden Markov Model. Tanabe and Wilbur (21) achieved 85.7% precision and 66.7% recall using a combination of statistical and knowledge-based strategies. They used a transformation-based part-of-speech tagger to recognize single-word protein names, and hand-crafted rules to filter out false positives and recover false negatives. Since the evaluation corpora used in these experiments were different from the corpus used in this paper, the results are not directly comparable.

Lee et al. (22) reported an F-measure of 69.2% on the GENIA corpus version 3.0 using SVMs. Shen et al. achieved an F-measure of 70.8% on the same corpus by incorporating various features into a Hidden Markov Model. Since the difference between the GENIA corpora versions 3.0 and 3.02, which we used in this paper, is small, their results suggest that their methods worked better than ours regarding recognition. However, their approaches do not provide ID information on recognized terms.

Krauthammer et al. (23) proposed a dictionary-based method of gene/protein name recognition. They used BLAST for approximate string matching by mapping sequences of text characters into sequences of nucleotides that could be processed by BLAST. They achieved a recall of 78.8% and a precision of 71.1% evaluated with a partial match criterion, which was not as stringent as our criterion.

8 Conclusion

We proposed a method of two-phase protein name recognition. In the first phase, we scan texts for protein name candidates using a protein name dictionary. In the second phase, we filter the candidates with machine learning. Our method is dictionary-based and can provide ID information on recognized terms, unlike machine learning approaches.

We presented two approaches to alleviating the low-recall problem caused by spelling variations. The first is to use an approximate string-searching algorithm instead of exact-matching algorithms. The second is to expand the dictionary in advance with the variant generator. We found the dictionary expansion approach to be much more attractive. The main reason was the cost of computation: the computational cost involved with approximate string searching was far bigger than for exact matching. Since there are huge amounts of available biomedical documentation, processing speed is an important factor in information extraction systems.

Experimental results using the GENIA corpus revealed that filtering using a naive Bayes classifier greatly improved precision with only a slight loss of recall, resulting in 10.8% improvement in F-measure, and dictionary expansion with the variant generator gave further 1.6% improvement and achieved an F-measure of 66.6%.

The future direction of this research involves:

- Use of state-of-the-art classifiers

We have used a naive Bayes classifier in our experiments because it re-

quires limited computational resources and performs well. There is a chance, however, of improving performance by using state-of-the-art machine learning techniques including maximum entropy models and support vector machines.

- Extending the algorithm for variant generation

Three types of operations were considered in this paper for the mechanism responsible for generating variants. There can, however, be other types of operations, such as word-insertion and word-replacement. Our future work should encompass these types of operations to improve the recall for long protein names.

- Applying the variant generator to other classes

We conducted experiments only on protein names in this paper. Since the variant generator can be easily applied to other classes, it would be interesting to investigate the performance gain on other classes. We expect to achieve comparable performance gain on at least gene and RNA names.

References

- [1] E. M. Marcotte, I. Xenarios, D. Eisenberg, Mining literature for protein-protein interactions, *BIOINFORMATICS* 17 (4) (2001) 359–363.
- [2] J. Thomas, D. Milward, C. Ouzounis, S. Pulman, M. Carroll, Automatic extraction of protein interactions from scientific abstracts, in: *Proceedings of the Pacific Symposium on Biocomputing (PSB2000)*, Vol. 5, 2000, pp. 502–513.
- [3] T. Ono, H. Hishigaki, A. Tanigami, T. Takagi, Automated extraction of information on protein-protein interactions from the biological literature, *BIOINFORMATICS* 17 (2) (2001) 155–161.

- [4] V. N. Vapnik, *The Nature of Statistical Learning Theory*, Springer, New York, 1995.
- [5] A. L. Berger, S. A. D. Pietra, V. J. D. Pietra, A maximum entropy approach to natural language processing, *Computational Linguistics* 22 (1) (1996) 39–71.
- [6] T. Ohta, Y. Tateisi, J.-D. Kim, J. Tsujii, Genia corpus: an annotated research abstract corpus in molecular biology domain, in: *Proceedings of the Human Language Technology Conference (HLT 2002)*, 2002.
- [7] K. Takeuchi, N. Collier, Use of support vector machines in extended named entity recognition, in: *Proceedings of the 6th Conference on Natural Language Learning 2002 (CoNLL-2002)*, 2002, pp. 119–125.
- [8] J. D. Kim, J. Tsujii, Corpus-based approach to biological entity recognition, in: *Text Data Mining SIG (ISMB2002)*, 2002.
- [9] J. Kazama, T. Makino, Y. Ohta, J. Tsujii, Tuning support vector machines for biomedical named entity recognition, in: *Proceedings of the ACL-02 Workshop on Natural Language Processing in the Biomedical Domain*, 2002, pp. 1–8.
- [10] D. Jurafsky, J. H. Martin, *Speech and Language processing*, Prentice-Hall, 2000.
- [11] G. Navarro, A guided tour to approximate string matching, *ACM Computing Surveys* 33 (1) (2001) 31–88.
- [12] E. S. Ristad, P. N. Yianilos, Learning string-edit distance, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20 (5) (1998) 522–532.
- [13] G. Navarro, R. Baeza-Yates, J. Arcoverde, Matchsimile: A flexible approximate matching tool for personal names searching, in: *Proceedings of the XVI Brazilian Symposium on Databases (SBBD'2001)*, 2001, pp. 228–242.

- [14] B. L. Humphreys, D. A. B. Lindberg, Building the unified medical language system, in: Proceedings of the 13th SCAMC, 1989, pp. 475–480.
- [15] D. D. Lewis, Naive Bayes at forty: The independence assumption in information retrieval., in: Proceedings of ECML-98, 10th European Conference on Machine Learning, no. 1398, 1998, pp. 4–15.
- [16] G. Escudero, L. Marquez, G. Rigau, Naive Bayes and exemplar-based approaches to word sense disambiguation revisited, in: Proceedings of the 14th European Conference on Artificial Intelligence, 2000, pp. 421–425.
- [17] T. Pedersen, A simple approach to building ensembles of naive Bayesian classifiers for word sense disambiguation, in: Proceedings of the First Annual Meeting of the North American Chapter of the Association for Computational Linguistics, 2000, pp. 63–69.
- [18] K. Nigam, R. Ghani, Analyzing the effectiveness and applicability of co-training, in: CIKM, 2000, pp. 86–93.
- [19] A. McCallum, K. Nigam, A comparison of event models for naive Bayes text classification, in: AAAI-98 Workshop on Learning for Text Categorization, 1998.
- [20] N. Collier, C. Nobata, J. Tsujii, Automatic acquisition and classification of molecular biology terminology using a tagged corpus, *Journal of Terminology* 7 (2) (2001) 239–258.
- [21] L. Tanabe, W. J. Wilbur, Tagging gene and protein names in biomedical text, *BIOINFORMATICS* 18 (8) (2002) 1124–1132.
- [22] K.-J. Lee, Y.-S. Hwang, H.-C. Rim, Two-phase biomedical NE recognition based on SVMs, in: Proceedings of the ACL 2003 Workshop on Natural Language Processing in Biomedicine, 2003, pp. 33–40.
- [23] M. Krauthammer, A. Rzhetsky, P. Morozov, C. Friedman, Using BLAST for identifying gene and protein names in journal articles, *Gene* 259 (2000)

245-252.